What Is the Carbon Premium a Premium On?*

Marcin Kacperczyk
Imperial College London

December 2025

Abstract

I develop a unified asset-pricing framework in which the priced climate characteristic is an equilibrium exposure function mapping firm fundamentals into cash-flow and discount-rate sensitivities. While a social-planner benchmark targets emission intensity through sustainable discount-rate wedges, market incompleteness, with heterogeneous climate concern and unspanned damages, shifts equilibrium pricing toward emission levels. I prove that intensity sufficiency is a restrictive knife-edge case. Moreover, endogenous technology can trigger a Jevons-type rebound in which efficiency gains increase equilibrium climate risk premia, generating a technological sign reversal. By reconciling the roles of levels and intensities, the model provides a formal protocol for identifying the carbon premium and clarifies conflicting empirical evidence.

*I thank Patrick Bolton, Michela Verardo, and Jiang Wang for helpful 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. Yet even at a conceptual level there is no agreement on what this premium is a premium on. Some studies measure exposure using absolute emissions (e.g., Bolton and Kacperczyk (2021, 2023)), others use emissions per unit of firm value (Pedersen (2025)), still others rely on changes in emissions, sectoral classifications (Pastor et al. (2022)), or more elaborate transition-risk scores (e.g., Sautner et al. (2023), Eskildsen et al. (2024)). At the same time, the policy debate over sustainable discount rates often proceeds as if there were a unique, well-defined wedge that should be applied to brown firms, typically expressed as a function of emission intensity. This combination of widespread use and conceptual ambiguity raises a basic question: what is the relevant exposure for a carbon premium in equilibrium asset prices?

A natural starting point is the complete-markets planner benchmark underlying the recent sustainable-discount-rate literature (e.g., Pedersen (2025)). In an economy with a representative agent, a well-defined social cost of carbon (SCC), and a full span of climate risk, a social planner chooses consumption and emissions paths to maximize discounted welfare. A standard result implies that the planner's allocation can be decentralized either with Pigouvian carbon taxes or with firm-specific sustainable discount rates that tilt the cost of capital in proportion to emission intensity. In particular, if cash-flow impacts are proportional to each firm's emission level $X_{i,t}$ and values are $V_{i,t}$ along the planner's allocation, then to a first-order approximation the deterministic wedge between financial and sustainable discount rates is proportional to $X_{i,t}/V_{i,t}$ times the gap between the SCC and the prevailing carbon price. In this benchmark, emission intensity is a sufficient statistic for how strongly a firm's discount rate must be adjusted to implement the planner's allocation, and it is tempting to conclude from this that "only X/V matters" for the greenium.

This paper shows that this conclusion is much less robust than it might appear. Once we depart from the complete-markets representative-agent benchmark and allow for incomplete risk sharing, heterogeneous climate concern, and richer technology and demand channels, equilibrium risk premia load on firms' climate exposures $g(Z_{i,t})$, which generally need not collapse to emissions intensity and may depend on emission levels and other firm characteristics. In incomplete markets, the mapping from the SCC to discount-rate wedges breaks down. The relevant carbon price of risk is an equilibrium object determined by the joint distribution of risk aversion and climate concern, and premia load on whatever characteristics govern firms' exposure to climate and policy shocks.

Intensity X/V emerges as a sufficient statistic only under a knife-edge set of conditions. Relaxing any one of these conditions yields pricing in emissions levels X (possibly channel-weighted) and/or in additional characteristics capturing demand, technology, or policy exposure.

Formally, I develop a tractable incomplete-markets model in which heterogeneous investors trade securities issued by emitting firms, face unhedgeable climate damages, and differ in their concern about those damages. I start with a framework in which firms' dividends load on a climate or policy shock in proportion to their emission levels X_i. In this environment there exists a unique aggregate stochastic discount factor that prices all traded assets, but it does not coincide with the intertemporal marginal rate of substitution of any social planner who internalizes the externality. The climate price of risk is instead determined in equilibrium by the cross-section of risk aversion and climate concern and by aggregate emissions. Under CARA-normal assumptions and their dynamic counterparts, expected excess returns are linear in emission levels, with

$$E[R_i] - R_f = \lambda X_i$$

in the static case and an analogous conditional relation in the dynamic case. The coefficient λ admits a transparent decomposition into a purely financial brown premium for bearing unhedgeable climate risk and a true greenium component reflecting the presence of investors who are willing to accept lower financial returns in order to reduce their exposure to climate damages. In contrast to the planner benchmark, the firm-specific exposure is the level of emissions X_i rather than the ratio X_i/V_i, and the equilibrium price of climate risk is not tied mechanically to the SCC.

I then relate this incomplete-markets benchmark to the complete-markets sustainable-discount-rate framework. Starting from a planner with a well-defined SCC and optimal Pigouvian taxation, I derive the familiar representation in which firm-specific sustainable discount rates are adjusted in proportion to $X_{i,t}/V_{i,t}$. This clarifies that the intensity-based wedge is a statement about implementing a first-best allocation under complete risk sharing, not a statement about the stochastic pricing of climate and transition risk in equilibrium. Once markets are incomplete and climate concern is heterogeneous, the same damage function that underpins the SCC generically leads to a level-based emission factor in the cross-section of expected returns, and the planner's intensity-based wedge and the market's level-based wedge diverge.

To move beyond these two benchmarks, I embed both environments in a general linear-factor setting for transition and climate risk. Dividends can load on one or more climate-related shocks

through exposure coefficients $g_k(Z_{i,t})$ that depend on a vector of firm characteristics $Z_{i,t}$, including, but not limited to, emissions $X_{i,t}$ and values $V_{i,t}$. A standard log-linear approximation to the Euler equation then implies that expected excess returns take the form

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

regardless of whether markets are complete or incomplete. Within this framework I characterize conditions under which emission intensity $X_{i,t}/V_{i,t}$ is a sufficient statistic for the pricing of a given climate or transition factor, in the sense that expected excess returns depend on firm i only through its intensity. I show that intensity is a unique sufficient statistic only if exposure functions satisfy restrictive conditions. Firms with the same intensity must have the same exposure, up to a common scalar, and the exposure must be homogeneous of degree one in $(X_{i,t}, V_{i,t})$. These conditions are naturally satisfied in the complete-markets planner benchmark with proportional cash-flow impacts per dollar of value and homogeneous pass-through, but they are generically violated once exposures depend on levels, changes, efficiency choices, or demand shifters.

Finally, I introduce a technological challenge to the intensity-sufficiency rule by embedding the model in a dynamic production setting with endogenous energy efficiency. I show that when firms can adjust their efficiency parameters, a Jevons-type rebound effect can emerge, whereby higher efficiency endogenously increases energy demand and emissions. Crucially, this mechanism can trigger a sign-reversal in risk sensitivity: improvements in efficiency can actually raise a firm's equilibrium risk premium by increasing its effective exposure to priced climate or policy shocks. This result demonstrates that emission intensity is not only an insufficient statistic but can be a misleading one. It fails to account for the technological capacity of firms to rebound into higher risk profiles. This establishes that the collapse of the carbon premium into a simple intensity metric is not just hindered by market frictions, but is fundamentally at odds with the endogenous nature of industrial technology.

Taken together, the analysis yields a sharp implication: the appropriate emissions measure depends on the economic object being priced. Emissions intensity $X_{i,t}/V_{i,t}$ is a natural sufficient statistic for *planner wedges* in complete-markets environments in which policy operates through a deterministic distortion to discounting. In contrast, equilibrium risk premia generically depend on the priced exposure functions $g_k(Z_{i,t})$ that map firm characteristics and policy-relevant states into cash-flow and discount-rate sensitivities. Once these exposures vary with technology, scale,

or other firm fundamentals, neither levels nor intensity is mechanically superior. Building on this structure, I provide testable sufficiency and spanning restrictions that turn the choice among levels, intensity, changes, and related metrics into a falsifiable empirical question. A candidate measure is valid only if it spans the return variation associated with the underlying exposure function. The framework therefore explains why competing measures can appear in the data, and it delivers an implementable protocol for adjudicating among them while keeping the economic mechanism, and hence the interpretation of climate-related pricing, transparent.

The framework delivers sharp falsifiable implications for empirical work. For example, if emission intensity is a uniquely sufficient statistic for climate-related risk, then expected returns must vary across firms only through $X_{i,t}/V_{i,t}$. Conversely, systematic pricing of emission levels or changes conditional on intensity directly rejects intensity sufficiency and implies that the priced climate exposure depends on a richer set of firm characteristics. To illustrate these distinctions empirically, I examine return regressions using an international panel of firms with scope 1 emissions. I estimate return-forecasting regressions in which one- and twelve-month-ahead excess stock returns are projected on log emissions, emission intensity, and changes in emissions, controlling for standard characteristics and fixed effects. The estimates show that absolute emissions are positively and significantly related to future returns, and that this relation remains robust when emission intensity is included, whereas intensity has little incremental explanatory power once levels are controlled for. Moreover, changes in emissions are rewarded at short horizons but not at annual horizons, consistent with a transition-risk exposure that depends on both levels and recent changes of emissions. These patterns are difficult to reconcile with a strict sustainable-discount-rate interpretation in which intensity is the unique sufficient statistic, but they arise naturally in an incomplete-markets equilibrium with heterogeneous climate exposure. Overall, the empirical evidence suggests that the restrictive homogeneity conditions required for intensity-sufficiency (as derived in Section 5) are rejected by the data. Instead, the results support a more general exposure function in which absolute scale (levels) remains a primary risk factor.

Related Literature. This paper relates to several strands of the climate-finance and sustainable-investing literature. On the theoretical side, my analysis is closely related to equilibrium models of sustainable investing 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 a greenium in expected returns. Pedersen et al. (2021) construct an ESG-

efficient frontier and an ESG-adjusted CAPM in a mean-variance setting with heterogeneously ESG-motivated investors, and show how ESG scores, including carbon intensity, affect equilibrium Sharpe ratios and expected returns. Zerbib (2022) and De Angelis et al. (2023) analyze sustainable investing when some investors internalize environmental externalities, deriving taste and exclusion premia in an S-CAPM and a dynamic climate-impact-investing model in which green investors' portfolio tilts raise brown firms' cost of capital and induce lower optimal emissions. Dangl et al. (2025) develop a unified equilibrium framework with deontological, non-consequentialist, and consequentialist social preferences and study how different preference types affect technology supply, aggregate harm, and cost-of-capital differences between green and brown firms, showing that social preferences can have substantial real effects even when observed cost-of-capital gaps are small and that such gaps are in general a poor proxy for real impact. These papers share the idea that heterogeneous ESG or climate preferences generate priced premia, but they typically take the relevant ESG or environmental metric as given. My framework instead provides a micro-foundation for these types by deriving them from equilibrium exposure functions.

A separate strand of research develops sustainable-discount-rate and social-cost-of-carbon frameworks and shows how Pigouvian taxes or cost-of-capital wedges can decentralize a planner's allocation. Pedersen (2025) shows that, in a complete-markets planner world with a given social cost of carbon, the sustainable-discount-rate wedge that implements the social optimum is proportional to firms' emission intensity. Chittaro et al. (2025) analyze when a carbon tax can be replicated by capital markets and show that, in a frictionless representative-agent economy, a carbon tax is equivalent to a return schedule that increases linearly in firms' emission intensity. Empirically, Pastor et al. (2025a) introduce the notion of a firm's carbon burden, defined as the present value of the social costs of its future emissions computed from social-cost-of-carbon paths and emission forecasts, and show that firms with higher carbon burdens have higher costs of capital even after controlling for past emissions. Their SCC-based burden is a forward-looking, level-based measure of the firm's carbon externality that is empirically priced. In my framework, such a burden corresponds to a particular exposure function $g(Z_i)$ built from emissions and SCC, and my analysis clarifies when SCC-based objects such as carbon burden can or cannot be interpreted as the relevant priced characteristic for the carbon premium in equilibrium asset prices.

On the empirical side, a large body of work documents how carbon emissions and greenness correlate with expected returns and the cost of capital, including evidence on carbon premia and screening by institutional investors in Bolton and Kacperczyk (2021, 2023) and broader green-

versus-brown return patterns in Pastor et al. (2022) and Eskildsen et al. (2024). More broadly, Pastor et al. (2025b) survey the rapidly growing sustainable-investing literature, emphasize the distinction between risk-based and taste-based components of green premia, and highlight as an open question which ESG attributes are actually priced. I address this question for carbon by asking, in a model with explicit emissions and climate risk, which emission-based characteristics can serve as sufficient statistics for exposure to the priced climate factors.

The remainder of the paper is organized as follows. Section 2 develops a simple static incomplete-markets model with heterogeneous investors and unspanned climate damages and derives a level-based emission factor together with a decomposition of the climate price of risk into a purely financial brown premium and a true greenium reflecting non-pecuniary climate concern. Section 3 then presents the complete-markets social-planner benchmark and the associated sustainable-discount-rate representation in which emission intensity summarizes the deterministic wedge between financial and sustainable discount rates. Section 4 embeds the incomplete-markets structure in a linear-factor environment with transition and climate shocks, while Section 5 characterizes the conditions under which emission intensity can be treated as a sufficient statistic for transition premia and shows how restrictive these conditions are once exposures depend on richer firm characteristics. Section 6 further illustrates the model dependence of emission-based risk premia by showing how different technologies and policy channels naturally generate premia on emission levels, intensities, and changes, and by providing a tractable specification in which both levels and changes in emissions emerge as priced characteristics. Section 7 extends the analysis to an environment with Jevons-type efficiency effects, highlighting how improvements in energy efficiency interact with emissions in shaping equilibrium risk premia and further undermining the uniqueness of intensity as a sufficient statistic. Finally, Section 8 discusses welfare and policy implications in light of the planner benchmark, derives empirical specifications that map the factor structure into cross-sectional return regressions using level-, intensity-, and change-based emission measures, and provides an illustrative empirical application using international firm-level data.

2 A Simple Static Model with Heterogeneous Investors and Climate Risk

Consider a two-period economy, with $t = 0, 1$. Uncertainty is represented by a finite state 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

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

I assume ζ is normally distributed and that 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 we further relax in Online Appendix C.

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 can be interpreted as a per-unit climate or policy shock that scales the cash-flow impact of emissions. For example, if ζ is a shock to the per-unit carbon cost or to the shadow value of emissions in production, then, for given X_i, a one-unit increase in ζ changes firm i's cash flow proportionally to X_i. Thus, conditional on the realization of ζ, the exposure of firm i to climate

risk is proportional to its emission *level* X_i.

Let P_i denote the date-0 price of firm i's equity, and define the gross return

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

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 stocks and a position b^h in the risk-free asset, subject to the budget constraint

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

Date-1 consumption of household h in state ω is

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

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

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

where $\alpha_h > 0$ is the coefficient of absolute risk aversion and $\phi_h \geq 0$ is the strength of climate concern for household h. The term $\phi_h \Gamma(\zeta)$ captures the disutility that household h derives from aggregate climate damages $\Gamma(\zeta)$, with $\phi_h = 0$ corresponding to a purely financial (brown) investor and $\phi_h > 0$ corresponding to a climate-conscious (green) investor.

In this environment, investors differ in their exposure to climate damages and trade only a limited set of assets. As a result, climate damages are not fully spanned in equilibrium, and discounting reflects an aggregation of heterogeneous investor marginal rates of substitution rather than the social planner's sustainable discount factor. The corresponding priced climate exposure is therefore an equilibrium object, generically a function of firm characteristics, whose associated premium can be decomposed into a risk-based brown component and a preference-driven true greenium, as formalized in Section 4.1.

3 Complete-markets Benchmark

The static incomplete-markets framework above provides a positive characterization of equilibrium asset pricing when climate damages enter investors' utility and are only partially internalized through heterogeneous portfolio choices. To assess the role of emissions from a normative perspective, it is useful to contrast this environment with a complete-markets benchmark in which a social cost of carbon (SCC) is well defined and a social planner internalizes the externality. This section restates, in a parsimonious setting, the planner logic underlying the sustainable-discount-rate framework and shows how the SCC translates into firm-specific discount-rate wedges that are proportional to emission intensity X_i/V_i. For transparency, I abstract from cross-sectional heterogeneity and financial risk and focus on deterministic cash flows and a representative decision maker.

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

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

$$U = u(C_0) + \beta\Big(u(C_1) - d(X_1)\Big), \tag{3.1}$$

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

$$X_1 = \sum_{i=1}^{N} X_i I_i. \tag{3.2}$$

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

$$C_0 + \sum_{i=1}^{N} I_i \leq G_0, \qquad C_1 = \sum_{i=1}^{N} Y_i I_i. \tag{3.3}$$

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 (3.1) subject to (3.2)-(3.3). 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

$$u'(C_0) = \lambda_0, \qquad \beta u'(C_1) = \lambda_1. \tag{3.4}$$

The first-order condition for I_i is

$$-\lambda_0 + \lambda_1 Y_i - \beta d'(X_1) X_i = 0. \tag{3.5}$$

Combining (3.4) and (3.5) yields

$$\frac{\lambda_0}{\lambda_1} = Y_i - \frac{\beta d'(X_1)}{\lambda_1} X_i = Y_i - \frac{d'(X_1)}{u'(C_1)} X_i. \tag{3.6}$$

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

$$S \equiv \frac{d'(X_1)}{u'(C_1)}. \tag{3.7}$$

Expression (3.6) can then be written as

$$\frac{\lambda_0}{\lambda_1} = Y_i - S X_i. \tag{3.8}$$

Equation (3.8) 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

$$q \equiv \frac{\lambda_1}{\lambda_0} = \beta \frac{u'(C_1)}{u'(C_0)}. \tag{3.9}$$

Then (3.8) is equivalent to

$$1 = q(Y_i - SX_i). \tag{3.10}$$

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

$$q \equiv \frac{\lambda_1}{\lambda_0} = \frac{1}{1+r}. \tag{3.11}$$

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

$$V_i = q(Y_i - \tau X_i) = \frac{Y_i - \tau X_i}{1+r}. \tag{3.12}$$

If $\tau = S$, then (3.10) and (3.12) coincide, implying that the competitive equilibrium with tax $\tau = S$ reproduces the first-best allocation.

3.2 Sustainable discount rates and the emission-intensity wedge

The sustainable-discount-rate approach asks whether, instead of imposing a tax equal to the SCC, one can implement the same allocation by suitably adjusting firms' discount rates while leaving their cash flows untaxed. In the one-period environment above, this logic 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 (3.12). 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,

$$\frac{Y_i - \tau X_i}{1+r} = \frac{Y_i}{1+r_i^x}. \tag{3.13}$$

Solving (3.13) for r_i^x gives

$$1 + r_i^x = \frac{Y_i}{Y_i - \tau X_i}(1+r). \tag{3.14}$$

Equation (3.14) expresses the sustainable discount rate as a nonlinear function of the financial rate

r, the tax τ, and the ratio $\tau X_i/Y_i$.

For small wedges and moderate discount rates, a first-order approximation in $\tau X_i/Y_i$ is informative. Writing $\varepsilon_i \equiv \tau X_i/Y_i$ and expanding the right-hand side of (3.14) to first order in ε_i yields

$$1 + r_i^x \approx (1 + r)(1 + \varepsilon_i),$$

so that

$$r_i^x - r \approx \varepsilon_i = \tau \frac{X_i}{Y_i}. \tag{3.15}$$

Let V_i denote the firm's value under the tax as in (3.12). For small τ, we have $V_i \approx Y_i/(1+r)$, and hence

$$\frac{X_i}{Y_i} \approx \frac{X_i}{(1+r)V_i} \approx \frac{X_i}{V_i},$$

where the last approximation absorbs the factor $1 + r$ into the first-order approximation (and is numerically innocuous at annual rates). Substituting into (3.15) yields the familiar emission-intensity wedge,

$$r_i^x \approx r + \tau \frac{X_i}{V_i}. \tag{3.16}$$

If the tax τ coincides with the SCC S in (3.7), then (3.16) links the sustainable discount rate directly to the SCC and emission intensity:

$$r_i^x \approx r + S \frac{X_i}{V_i}. \tag{3.17}$$

More generally, suppose that the prevailing carbon tax is τ while the planner's SCC is S. The additional wedge required to move from the current tax to the first-best tax is $S - \tau$. Consider a small change $\mathrm{d}\tau$ in the tax and equate the induced change in firm value to that from a small change $\mathrm{d}r_i$ in the discount rate:

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

Setting $\mathrm{d}V_i^{(\tau)} = \mathrm{d}V_i^{(r)}$ implies

$$\mathrm{d}r_i = \frac{X_i}{V_i}\mathrm{d}\tau.$$

Applying this with $\mathrm{d}\tau = S - \tau$ and linearizing around the initial allocation yields

$$r_i^x \approx r + (S - \tau)\frac{X_i}{V_i}. \tag{3.18}$$

Equations (3.7) and (3.18) formalize the complete-markets benchmark. The SCC S is the planner's shadow value of emissions in units of the consumption numeraire, and, to a first–order approximation, 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 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. Subsequent sections show that once one relaxes the complete-markets and representative-agent assumptions and allows for heterogeneous climate concern and incomplete risk sharing as in Section 2.3, the priced characteristics that govern equilibrium risk premia need not collapse to X_i/V_i and can instead depend on emission levels and more general exposure functions.

4 Equilibrium with Heterogeneous Investors and Unspanned Climate Risk

The benchmark models in the previous sections clarify two distinct objects: equilibrium risk premia in incomplete markets and implementation wedges in a complete-markets planner economy. What they do not yet provide is a general way to think about how climate and transition risk enter asset prices when firms are exposed through multiple, overlapping channels. In practice, emissions, technology, demand, and policy all shape how firms load on climate-related shocks, and no single emission measure need summarize these exposures a priori. This section introduces a general factor framework that makes these channels explicit and provides a unified language for assessing when emission intensity is sufficient for pricing and when it is not.

4.1 Incomplete markets

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

Markets are *incomplete* in that there is no portfolio of traded assets whose payoff replicates

the damage term $\Gamma(\zeta) = \kappa \bar{X} \zeta$. Equivalently, ζ is not fully spanned by the space of traded payoffs once the non-pecuniary climate term in preferences is taken into account. 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.

Formally, the equilibrium still admits a unique stochastic discount factor m that prices all traded assets, but the incomplete spanning of $\Gamma(\zeta)$ prevents us from representing m as the marginal rate of substitution of a representative planner who internalizes the externality in a first-best sense. This is exactly where the analysis differs from the complete-markets benchmark.

4.2 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) is

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

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

$$\mathbb{E}[C_1^h - \phi_h \Gamma(\zeta)] - \frac{\alpha_h}{2} \mathrm{Var}[C_1^h - \phi_h \Gamma(\zeta)]. \tag{4.1}$$

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

Using $\mathbb{E}[\zeta] = 0$ and $\mathrm{Var}(\zeta) = \sigma_\zeta^2$, objective (4.1) becomes

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

4.3 First-order conditions and equilibrium prices

Let $\boldsymbol{X} = (X_1, \ldots, X_N)^\top$ and write $\eta \boldsymbol{X}$ for the vector with components ηX_i. The first-order condition for θ_i^h is

$$\mu_i - P_i R_f - \alpha_h \sigma_\zeta^2 \Big(\sum_{j=1}^{N} \theta_j^h \eta X_j - \phi_h \kappa \bar{X} \Big) \eta X_i = 0, \qquad i = 1, \ldots, N. \tag{4.2}$$

Define the scalar

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

Then the first-order condition (4.2) can be written as

$$\mu_i - P_i R_f = \alpha_h \sigma_\zeta^2 \delta_h \eta X_i, \qquad i = 1, \ldots, N. \tag{4.3}$$

The left-hand side of the above equation does not depend on h, so in equilibrium the right-hand side must coincide across all households h that hold stock i.

Proposition 1 (One-factor structure and emission levels)**.** *In any competitive equilibrium with $\eta > 0$, there exists a scalar λ such that*

$$\mu_i - P_i R_f = \lambda \eta X_i, \qquad i = 1, \ldots, N, \tag{4.4}$$

and for every household h,

$$\delta_h = \frac{\lambda}{\alpha_h \sigma_\zeta^2}. \tag{4.5}$$

Proof. By (4.3), for any h that holds a positive amount of stock i we have

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

Since this ratio must be equalized across all stocks i and all households h with interior holdings, it

must equal a common scalar λ:

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

which yields (4.4). Substituting back gives $\alpha_h \sigma_\zeta^2 \delta_h = \lambda$, which yields (4.5). □

Equation (4.4) is a linear factor structure in the *level* of emissions X_i:

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

where $\mathbb{E}[R_i] = \mu_i / P_i$. The coefficient λ is the price of climate-related risk. Emission levels X_i are the relevant exposure.

4.4 Solving for the price of climate risk

I now solve explicitly for λ in terms of (α_h, ϕ_h) and the equilibrium emissions profile.

Multiplying the definition of δ_h by η and summing over h gives

$$\sum_{h=1}^{H} \sum_{j=1}^{N} \theta_j^h \eta^2 X_j - \kappa \bar{X} \eta \sum_{h=1}^{H} \phi_h = \eta \sum_{h=1}^{H} \delta_h. \tag{4.6}$$

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

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

By (4.5), the right-hand side of (4.6) is

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

Equating both sides and solving for λ yields:

Proposition 2 (Price of climate risk)**.** *In the static model above, the equilibrium price of climate risk is*

$$\lambda = \frac{\sigma_\zeta^2 \bar{X}}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \Big(\eta - \kappa \sum_{h=1}^{H} \phi_h \Big). \tag{4.7}$$

Under CARA preferences, λ depends on the distribution of risk aversion (α_h) and the aggregate climate concern $\sum_h \phi_h$ and it is independent of initial wealth (W_0^h). We show in Online Appendix B that with CRRA utility the analogous expression would feature wealth weights.

4.5 Decomposing transition risk and preference effects

Expression (4.7) admits a natural decomposition into a purely financial brown risk premium (transition-risk component) and a true greenium associated with non-pecuniary climate concern (preference component). To make this explicit, note the price of climate risk in the incomplete-markets economy derived above

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

Define

$$\lambda_{X,t}^{\mathrm{TR}} \equiv \frac{\sigma_\zeta^2 \bar{X}_t \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}}, \qquad \lambda_{X,t}^{\mathrm{P}} \equiv -\frac{\sigma_\zeta^2 \bar{X}_t \kappa \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \sum_{h=1}^{H} \phi_h, \tag{4.8}$$

so that

$$\lambda_{X,t} = \lambda_{X,t}^{\mathrm{TR}} + \lambda_{X,t}^{\mathrm{P}}. \tag{4.9}$$

The term $\lambda_{X,t}^{\mathrm{TR}}$ is the price of climate risk that would prevail if all households were climate-neutral in their preferences, 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 $D_{i,t+1} = \mu_{i,t+1} + \eta X_{i,t} \zeta_{t+1}$. In this sense, $\lambda_{X,t}^{\mathrm{TR}}$ captures the transition-risk or physical-risk component of the brown premium. It depends on the volatility of the climate shock σ_ζ^2, aggregate emissions $\bar{X}_t$, the technological sensitivity η, and aggregate risk tolerance $\sum_h (1/\alpha_h)$, but it is independent of the climate-preference parameters $\{\phi_h\}_h$.

By contrast, the term $\lambda_{X,t}^{\mathrm{P}}$ is zero when $\phi_h = 0$ and is strictly decreasing in aggregate climate concern $\sum_h \phi_h$ whenever $\kappa > 0$. Differentiating (4.8) with respect to ϕ_h yields

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

so that an increase in the weight that household h puts on climate damages lowers the equilibrium price of climate risk. The component $\lambda_{X,t}^{\mathrm{P}}$ therefore captures the preference or true greenium effect: for a given technology and shock process, stronger climate concern tilts portfolios away from high-emission exposures and compresses premia on brown assets.

Given the decomposition (4.9), the expected excess return on firm i can be written as

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx (\lambda_{X,t}^{\mathrm{TR}} + \lambda_{X,t}^{\mathrm{P}}) X_{i,t} = \lambda_{X,t}^{\mathrm{TR}} X_{i,t} + \lambda_{X,t}^{\mathrm{P}} X_{i,t}. \tag{4.10}$$

In this decomposition, the first term is the compensation for bearing non-diversifiable transition (or physical) climate risk, as encoded in the distribution of ζ_{t+1} and its impact on cash flows. The second term reflects the fact that, in incomplete markets, investors who directly dislike damages are willing to accept lower expected returns on low-emission assets (or require higher returns on high-emission assets) even when the underlying shock process is held fixed.

An equivalent separation applies at the level of the aggregate stochastic discount factor. Writing the affine SDF as

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

and substituting (4.8) gives

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

The loading B_t^{TR} describes how the representative pricing kernel would respond to the climate shock in an otherwise identical economy with $\phi_h = 0$. The adjustment B_t^{P} captures the additional tilt in the pricing kernel induced by heterogeneous climate preferences under incomplete markets.

This decomposition makes it clear that, within our framework, transition risk and preference effects are conceptually distinct but jointly determine the observed carbon premium. The former reflects the technological and policy sensitivity of cash flows to ζ_{t+1}, while the latter reflects investors' willingness to trade off financial returns against exposure to climate damages.

4.6 When is emission intensity a unique sufficient statistic?

The previous subsection shows that, in a general linear-factor environment, the cross-section of expected excess returns is governed by firms' factor exposures, which are allowed to depend on a rich set of characteristics $Z_{i,t}$. There is therefore no a priori reason why any single primitive, such as emission levels $X_{i,t}$ or intensity $X_{i,t}/V_{i,t}$, must be sufficient to summarize climate-related risk. In the baseline, the logic is naturally level-based. When emissions enter through firms' cash-flow or discount-rate sensitivities, expected excess returns load on $X_{i,t}$, motivating level-based specifications. This subsection characterizes the conditions under which emission intensity is, in fact, a unique sufficient statistic for pricing a given climate or transition factor, in the sense that expected excess returns depend on firm i only through $X_{i,t}/V_{i,t}$.

I begin with a formal definition.

Definition 4.1 (Emission intensity as a sufficient statistic)**.** *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 i's 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, in the sense that there exists a scalar ψ_t such that

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

We work within a general linear-factor framework. For ease of notation, we suppress the time index t when no confusion can arise.

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

$$D_{i,t+1} = \mu_{i,t+1} + g(Z_{i,t}) f_{t+1}, \tag{4.11}$$

where $Z_{i,t}$ collects firm characteristics (including in particular $(X_{i,t}, V_{i,t})$), and $g(\cdot)$ is a deterministic function. Let the stochastic discount factor be

$$m_{t+1} = a_t + b_t f_{t+1} + \varepsilon_{t+1}, \tag{4.12}$$

with $a_t > 0$, b_t $\mathcal{F}_t$–measurable, and ε_{t+1} conditionally mean-zero and orthogonal to f_{t+1}. A standard log-linear approximation to the Euler equation implies that, to a first-order approximation in the conditional variance of f_{t+1}, expected excess returns take the form

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx \lambda_t g(Z_{i,t}), \qquad \lambda_t = -\frac{b_t \mathrm{Var}_t(f_{t+1})}{\mathbb{E}_t[m_{t+1}]}, \tag{4.13}$$

regardless of whether markets are complete or incomplete. The role of market structure is to

determine b_t and hence λ_t, but not the functional form of $g(\cdot)$.

The next result provides conditions under which $g(Z_{i,t})$ depends on firms only through the emission intensity $X_{i,t}/V_{i,t}$ and, moreover, under which this dependence is unique up to scale.

Proposition 3 (Conditions for emission intensity to be a unique sufficient statistic)**.** *Consider the one-factor environment* (4.11)-(4.13) *at a given date* t*. Suppose that:*

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

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

contains an open subset of $\mathbb{R}^2_+$ *(when viewed in the limit as* $N \to \infty$*).*

(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*,*

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

where we explicitly highlight the dependence of g *on* $(X_{i,t}, V_{i,t})$*.*

(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 exposures coincide,*

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

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

$$g(Z_{i,t}) = \varphi_t\left(\frac{X_{i,t}}{V_{i,t}} \right) \quad \text{for all } i. \tag{4.14}$$

If, in addition, $g(\cdot)$ *is continuous on* $\mathcal{S}_t$ *and the mapping* $v \mapsto \varphi_t(u)$ *is linear in* v *for each fixed* u*, then there exists a scalar* κ_t *such that*

$$g(Z_{i,t}) = \kappa_t \frac{X_{i,t}}{V_{i,t}} \quad \text{for all } i, \tag{4.15}$$

and $X_{i,t}/V_{i,t}$ *is a unique sufficient statistic for expected excess returns in the sense of Definition 4.1.*

Proof. 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 $g(Z_{i,t})$. Define $\varphi_t(u)$ to be this common value. 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 (4.14).

To obtain (4.15), consider 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 $\lambda > 0$

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

The left-hand side, by condition (iii), depends only on the intensity u (which is invariant to the scaling) and on $\lambda V_{i,t}$. Writing the left-hand side as $\varphi_t(u; \lambda V_{i,t})$ and the right-hand side as $\lambda \varphi_t(u; V_{i,t})$, we obtain

$$\varphi_t(u; \lambda V_{i,t}) = \lambda \varphi_t(u; V_{i,t}).$$

Under the stated linearity assumption in v, this functional equation implies that, for each fixed u, $\varphi_t(u; v)$ must be of the form $\kappa_t(u)v$ for some scalar $\kappa_t(u)$. Since intensity $u = X_{i,t}/V_{i,t}$ and size $v = V_{i,t}$ jointly determine $(X_{i,t}, V_{i,t})$, and $g(Z_{i,t})$ is homogeneous of degree one, it follows that

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

If, in addition, $\kappa_t(u)$ is constant in u on the relevant support, we obtain (4.15) with κ_t independent of i. In that case, any alternative characteristic $\tilde{\chi}_{i,t}$ that yields the same one-factor pricing relation (4.13) must satisfy $\tilde{\chi}_{i,t} = c_t g(Z_{i,t})$ for some nonzero scalar c_t, and hence $\tilde{\chi}_{i,t}$ is proportional to $X_{i,t}/V_{i,t}$, establishing uniqueness in the sense of Definition 4.1. □

By symmetry, while Proposition 3 defines the restrictive homogeneity conditions required for intensity-sufficiency, the pricing of absolute levels emerges as the unique priced characteristic whenever the exposure function is invariant to firm value ($\frac{\partial g}{\partial V} = 0$). This property is naturally satisfied by most linear policy instruments, such as a carbon tax per unit of emissions, whereas the homogeneity required for intensity requires a specific, non-linear interaction between emissions and market capitalization.

Proposition 3 encompasses both complete- and incomplete-markets environments. The conditions are imposed on the exposure function $g(\cdot)$, which is determined by technology and the way the factor enters dividends, and not on the SDF. Market completeness affects the determination of the factor price of risk λ_t in (4.13) (through the coefficients a_t and b_t in (4.12)), but it does not alter the functional dependence of $g(Z_{i,t})$ on $(X_{i,t}, V_{i,t})$.

In the complete-markets planner benchmark of Section 3, with a deterministic SCC path $\{S_t\}_t$ and a tax channel in which cash-flow losses are proportional to $X_{i,t}$ per unit of emissions, the mapping from SCC to sustainable discount rates can be expressed in terms of $X^*_{i,t}/V^*_{i,t}$ as in (3.18). Under additional assumptions on technology and capital markets, this environment can be embedded in the class covered by Proposition 3, with $g(Z_{i,t})$ proportional to $X_{i,t}/V_{i,t}$ along the planner's allocation, so that intensity is a sufficient statistic for the *deterministic* wedge. By contrast, in the equilibrium environments developed below in Sections 5 and 6, the linear exposure coefficients $g_k(Z_{i,t})$ need not satisfy conditions (ii)-(iii). Once $g_k(\cdot)$ depends on additional characteristics (such as levels, changes, efficiency, or demand shifters), the approximation in (4.13) implies that expected excess returns need not be functions of intensity alone, and the uniqueness of $X_{i,t}/V_{i,t}$ as a sufficient statistic for transition-risk premia is lost.

In sum, the results in this section demonstrate that once we depart from the complete-markets planner benchmark, the pricing of carbon risk is no longer constrained by the normative requirement of intensity sufficiency. When markets are incomplete and damages are unspanned, the market equilibrium forces a divergence in that a planner may still target intensity to internalize externalities, but investors price the risk associated with the absolute scale of exposure. This tension suggests that the intensity-only rule is not a robust feature of climate asset pricing, but rather a specific outcome of the spanning assumptions. This raises a more fundamental question: under what formal mathematical conditions would intensity ever emerge as a sufficient statistic? To answer this, we must move beyond specific market frictions and characterize the general properties of the exposure function itself.

5 Transition Risk, Emission Intensity, and Priced Characteristics

Building on the intuition that market incompleteness breaks the intensity-sufficiency rule, this section provides a formal characterization of the priced climate characteristic in a general linear-factor environment. Rather than assuming a specific friction, I reverse-engineer the problem in that I define the climate exposure function $g(Z_{i,t})$ as the equilibrium mapping from firm fundamentals to risk sensitivities and derive the mathematical restrictions required for this mapping to collapse into the ratio $X_{i,t}/V_{i,t}$.

By abstracting from the specific mechanisms of Section 4, one can identify the knife-edge conditions–specifically regarding the homogeneity of the dividend and damage processes–under

which intensity is a sufficient statistic for the carbon premium. This general approach allows us to adjudicate between competing empirical measures and clarifies why, in a world of non-linear damages and evolving technology, the priced characteristic $g(Z)$ will generically depend on a richer set of attributes than intensity alone.

5.1 Sustainable discount rates and intensity in the complete-markets benchmark

In the complete-markets benchmark with a representative planner and a well-defined social cost of carbon $\{S_t\}_t$, a Pigouvian tax sequence $\{\tau_t\}_t$ or an equivalent cap-and-trade scheme decentralizes the planner's allocation. Let $V^*_{i,t}$ denote the value of firm i at date t along this allocation and let $X^*_{i,t}$ denote its emissions. For a given financial discount rate r_t, one can define firm-specific *sustainable* discount rates $r^x_{i,t}$ by equating the present value of taxed cash flows discounted at r_t with the present value of untaxed cash flows discounted at $r^x_{i,t}$. A first-order expansion around a baseline with zero wedge yields

$$r^x_{i,t} \approx r_t + (S_t - \tau_t)\frac{X^*_{i,t}}{V^*_{i,t}}, \tag{5.1}$$

where $S_t - \tau_t$ is the gap between the SCC and the prevailing carbon price. In this sense, under complete markets and optimal taxation, the *deterministic* wedge between financial and sustainable discount rates is summarized by the emission intensity $X^*_{i,t}/V^*_{i,t}$.

Crucially, (5.1) is a statement about the *level* of the discount-rate wedge required to implement the planner's allocation for given S_t, not a statement about the *pricing of stochastic climate or policy risk*. The intensity $X^*_{i,t}/V^*_{i,t}$ appears because cash-flow losses proportional to $X^*_{i,t}$ are being expressed on a per-unit-of-value basis.

5.2 Transition risk as stochastic exposure to policy or technology shocks

Consider now transition risk, by which we mean stochastic shocks to future climate policy, carbon prices, or technology that affect 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

$$D_{i,t+1} = \tilde{D}_{i,t+1} - \tau_{t+1}X_{i,t}, \tag{5.2}$$

where $\tilde{D}_{i,t+1}$ collects all other components orthogonal (conditionally) to τ_{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

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

with $\mu_{i,t+1}$ conditionally independent of $\hat{\tau}_{t+1}$. In any equilibrium with a strictly positive SDF m_{t+1} that is conditionally square-integrable, the expected excess return on asset i satisfies

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

Under a standard log-linear approximation, and assuming that $\hat{\tau}_{t+1}$ is a factor in m_{t+1}, this yields a linear factor structure in which the *transition-risk premium* associated with $\hat{\tau}_{t+1}$ takes the form

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx \lambda_t^{(\tau)} \theta_{i,t}^{(\tau)} = -\lambda_t^{(\tau)} X_{i,t}, \tag{5.3}$$

where the price of transition risk $\lambda_t^{(\tau)}$ is proportional to $-\mathrm{Cov}_t(m_{t+1}, \hat{\tau}_{t+1})$.

Expression (5.3) shows that, in this simple tax channel, the natural priced characteristic for transition risk is the *level* of emissions $X_{i,t}$. One can trivially rewrite $X_{i,t} = (X_{i,t}/V_{i,t})V_{i,t}$ and therefore express $\theta_{i,t}^{(\tau)}$ in terms of intensity and size, but there is no sense in which $X_{i,t}/V_{i,t}$ is uniquely singled out as the sufficient statistic. The exposure coefficient $\theta_{i,t}^{(\tau)}$ is determined by the way the shock $\hat{\tau}_{t+1}$ enters the cash flow, not by the representative agent's planner problem.

More generally, if we allow for multiple transition channels, including demand shocks, stranded-asset risk, and technology shocks, the dividend of firm i takes the generic form

$$D_{i,t+1} = \tilde{D}_{i,t+1} + \sum_{k=1}^{K} g_k(Z_{i,t}) f_{k,t+1}, \tag{5.4}$$

where $f_{k,t+1}$ are mean-zero transition or climate factors and $g_k(Z_{i,t})$ are exposure coefficients that depend on a vector of firm 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. In this representation, the priced characteristic associated with factor k is $g_k(Z_{i,t})$, not any particular primitive such as $X_{i,t}$ or $X_{i,t}/V_{i,t}$ taken in isolation.

5.3 Complete versus incomplete markets and the priced characteristic

The distinction between complete and incomplete markets affects the *price* of transition risk but not, at a given technology and shock structure, the functional form of the exposure coefficients $g_k(Z_{i,t})$. 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 the dividend of asset i can be written as in (5.4):

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

with $\mu_{i,t+1}$ conditionally independent of $\{f_{k,t+1}\}_k$ and $Z_{i,t}$ $\mathcal{F}_t$–measurable. Assume that there exists an SDF m_{t+1} of the form

$$m_{t+1} = a_t + \sum_{k=1}^{K} b_{k,t} f_{k,t+1} + \varepsilon_{t+1}, \tag{5.5}$$

where $a_t > 0$, $\{b_{k,t}\}_k$ are $\mathcal{F}_t$–measurable, and ε_{t+1} is a mean-zero residual orthogonal to all 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 in general, but in our CARA-normal environment a particular affine representation is pinned down by equilibrium.

Under a local linear approximation in which higher-order terms in the factors are negligible, the Euler equation $\mathbb{E}_t[m_{t+1}R_{i,t+1}] = 1$ implies the following characterization of expected excess returns.

Proposition 4 (Priced characteristics in a linear factor structure)**.** *Suppose that dividends are given by* (5.4) *and that the SDF admits the representation* (5.5)*, with* $\{f_{k,t+1}\}_k$ *mutually uncorrelated conditional on* $\mathcal{F}_t$*. Then, up to a first-order approximation in the conditional variances of the factors,*

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx \sum_{k=1}^{K} \lambda_{k,t} g_k(Z_{i,t}), \tag{5.6}$$

where

$$\lambda_{k,t} = -\frac{b_{k,t}\mathrm{Var}_t(f_{k,t+1})}{\mathbb{E}_t[m_{t+1}]}, \qquad k = 1, \dots, K, \tag{5.7}$$

are the factor prices of risk. In particular, for each factor k*, the* priced characteristic *is* $g_k(Z_{i,t})$*, that is, two assets with identical* $g_k(Z_{i,t})$ *for all* k *have the same expected excess return in this approximation, regardless of how* $g_k(\cdot)$ *depends on the primitives* $(X_{i,t}, V_{i,t}, \dots)$*.*

The proof is immediate from substituting (5.4) and (5.5) into the Euler equation, using condi-

tional mutual uncorrelatedness of the factors and discarding higher-order covariance terms. The result holds irrespective of whether markets are complete or incomplete; completeness affects the determination of the coefficients $\{b_{k,t}\}_k$ and thus the prices of risk $\{\lambda_{k,t}\}_k$, but not the functional form of the exposure coefficients $g_k(Z_{i,t})$, which are pinned down by technology and the way shocks enter cash flows.

In the specific transition-tax example (5.2) with a single factor $f_{1,t+1} = \hat{\tau}_{t+1}$, Proposition 4 identifies $g_1(Z_{i,t})$ with $-X_{i,t}$. Once we allow for more general specifications in which $g_k(Z_{i,t})$ may involve levels, intensities, changes, or interactions, there is no reason to expect that a single statistic such as $X_{i,t}/V_{i,t}$ will be sufficient to summarize transition-risk exposure across all models. In the next section, I show by example how alternative, economically motivated specifications of $g_k(\cdot)$ generate premia on emission levels, intensity, and changes, and how simple economic restrictions discipline the admissible forms of $g_k(Z_{i,t})$.

6 Illustrative Exposure Specifications

6.1 General setup and discipline of exposure functions

The incomplete-markets derivation above establishes that, under a particular specification of the dividend exposure to the climate shock, expected excess returns are proportional to the level of emissions $X_{i,t}$. It is important, however, not to over-interpret this special case as implying that $X_{i,t}$ is the uniquely relevant priced quantity. In this section, I use the general factor representation in Proposition 4 to illustrate how different modeling choices for the exposure coefficient $g_k(Z_{i,t})$ generate a variety of emission-based premia, and to discuss how economic structure disciplines which metrics are admissible.

Consider first a single transition or climate factor ζ_{t+1} and write firm i's dividend as

$$D_{i,t+1} = \tilde{D}_{i,t+1} + \theta_{i,t}\zeta_{t+1}, \tag{6.1}$$

where $\tilde{D}_{i,t+1}$ collects all components of the payoff that are orthogonal to ζ_{t+1} conditional on information at t. Let $Z_{i,t}$ denote a vector of firm characteristics at date t, which may include emissions, value, size, sector, and other state variables. Suppose that the exposure coefficient is given by a measurable function

$$\theta_{i,t} = g(Z_{i,t}), \tag{6.2}$$

for some $g : \mathcal{Z} \to \mathbb{R}$ defined on the set of admissible characteristics $\mathcal{Z}$. In the special case of Section 4 we have $Z_{i,t} = X_{i,t}$ and $g(z) = \eta z$, yielding $\theta_{i,t} = \eta X_{i,t}$. Under the assumptions of Proposition 4, the conditional pricing relation implied by (6.1)-(6.2) is

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t g(Z_{i,t}), \tag{6.3}$$

for some scalar price of climate risk λ_t determined by preferences, technology, and the distribution of aggregate exposures. Equation (6.3) makes it clear that the priced object in the cross section is $g(Z_{i,t})$, the exposure coefficient implied by the specification (6.2), rather than any particular component of $Z_{i,t}$.

Different specifications of g correspond to different economically meaningful forms of climate risk. If $Z_{i,t} = (X_{i,t}, V_{i,t})$ and $g(X_{i,t}, V_{i,t}) = \eta X_{i,t}/V_{i,t}$, then (6.3) implies

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t \eta \frac{X_{i,t}}{V_{i,t}},$$

so that emission intensity $X_{i,t}/V_{i,t}$ is the relevant exposure. If instead g is specified as a linear combination,

$$g(Z_{i,t}) = \eta_1 X_{i,t} + \eta_2 \frac{X_{i,t}}{V_{i,t}},$$

then

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t \eta_1 X_{i,t} + \lambda_t \eta_2 \frac{X_{i,t}}{V_{i,t}},$$

so that both levels and intensity enter expected returns as priced components of the same climate-risk factor. More generally, g may depend on changes in emissions, $Z_{i,t} = (X_{i,t}, X_{i,t} - X_{i,t-1})$, on industry or technology indicators, or on firm-specific policy risk indicators, thereby generating premia on growth rates of emissions, sector-specific emission profiles, or other combinations of environmental characteristics.

The same reasoning extends to a multi-factor structure. Suppose that climate and policy risk are represented by a K-dimensional vector of shocks

$$\boldsymbol{\zeta}_{t+1} = (\zeta_{1,t+1}, \ldots, \zeta_{K,t+1})^\top,$$

and that firm i's dividend exposure is

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

for functions $g_k : \mathcal{Z} \to \mathbb{R}$. Under the conditions of Proposition 4, the linearized pricing relation becomes

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \sum_{k=1}^{K} \lambda_t^{(k)} g_k(Z_{i,t}), \tag{6.4}$$

where $\lambda_t^{(k)}$ is the equilibrium price of the kth climate or policy risk. Each choice of $(g_1, \ldots, g_K)$ corresponds to a different mapping from primitive firm characteristics to priced exposures. Emission levels, intensities, changes, and interacted terms with firm value or sectoral indicators all arise as special cases of (6.4) under appropriate specifications of $(g_k)_{k=1}^K$.

In this general factor representation, the exposure coefficients $g_k(Z_{i,t})$ are not arbitrary regressors but are disciplined by economic structure. For a given climate or transition factor $\zeta_{k,t+1}$, $g_k(Z_{i,t})$ is the derivative of the present value of firm i's cash flows (and, in the planner benchmark, damages) with respect to $\zeta_{k,t+1}$, evaluated along a baseline path. Only firm characteristics that enter the dividend, technology, policy, or damage blocks can therefore appear in $Z_{i,t}$, and simple invariance restrictions further restrict their functional form: for example, splitting a firm into two identical copies with half the value and half the emissions should not change total exposure, and, holding technology fixed, a higher emission level should not reduce exposure to a positive damage shock. In this sense, candidate metrics such as emission levels, changes in emissions, or SCC-weighted carbon burden measures are admissible only if they can be derived from a coherent specification of how $\zeta_{k,t+1}$ affects firm cash flows and damages, and they are empirically sufficient in the sense that, conditional on $g_k(Z_{i,t})$, other observable firm characteristics do not earn a distinct premium for that factor.

6.2 Transition losses depending on levels and changes

In the baseline incomplete-markets environment, firms' dividends load on a transition or climate factor ζ_{t+1} through exposure coefficients $g_\zeta(Z_{i,t})$ that are functions of firm characteristics $Z_{i,t}$, and the conditional pricing relation takes the form

$$E_t[R_{i,t+1} - R_{f,t}] \;=\; \lambda_{\zeta,t}\, g_\zeta(Z_{i,t}),$$

where $\lambda_{\zeta,t}$ is the equilibrium price of transition risk. In the simplest version of the model we have $g_\zeta(Z_{i,t}) \propto X_{i,t}$, so that expected returns load on emission levels. This subsection shows how changes in emissions, in addition to levels, can emerge as natural priced characteristics by modifying only the transition-risk technology, leaving the stochastic discount factor and the incompleteness structure unchanged.

Consider a one-factor transition-risk block in which there is a common policy-tightening or transition shock $J_{t+1} \in \{0, 1\}$ with $E_t[J_{t+1}] = p_t$ and innovation

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

The shock J_{t+1} captures an unanticipated tightening of climate policy, an abrupt shift in the transition path, or a jump in the effective carbon price. Suppose that, conditional on J_{t+1}, firm i's dividend at $t+1$ can be written as

$$D_{i,t+1} \;=\; \bar{D}_{i,t+1} \;-\; \Gamma_{i,t}\, J_{t+1}, \tag{6.5}$$

where $\bar{D}_{i,t+1}$ is the dividend in the absence of a tightening shock and $\Gamma_{i,t}$ is the loss in dividend value if the tightening shock occurs. We interpret $\Gamma_{i,t}$ as the present value (at t) of additional abatement costs, stranded-asset write-downs, and other transition losses that are realized when policy tightening arrives one period earlier than expected.

To link $\Gamma_{i,t}$ to emissions, we assume that the size of the transition loss depends on both the firm's current emission level $X_{i,t}$ and the recent change in emissions

$$\Delta X_{i,t} \;=\; X_{i,t} - X_{i,t-1}.$$

The level $X_{i,t}$ captures how dirty the firm is at date t, while $\Delta X_{i,t}$ captures whether the firm has recently increased or decreased its reliance on high-emission technologies. A firm that is both highly emitting and still ramping up emissions is more exposed to an abrupt policy tightening than a firm that is equally dirty but rapidly decarbonizing. Formally, I assume that the transition loss takes the linear form

$$\Gamma_{i,t} \;=\; \phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}, \tag{6.6}$$

with $\phi_1 > 0$ and $\phi_2 \geq 0$. The coefficient ϕ_1 scales the loss with the level of emissions, while ϕ_2 captures the additional sensitivity of the loss to recent emission growth or reduction.

Substituting (6.6) into (6.5) and using the definition of ζ_{t+1}, we can write the dividend as

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

The first term in the second line is the expected loss under the current transition path, while the last term is the unexpected component associated with the innovation ζ_{t+1}. The loading of firm i's dividend on the transition factor is therefore

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

where we can treat $Z_{i,t}$ as containing $(X_{i,t}, \Delta X_{i,t})$ among other characteristics. The sign convention for g_ζ is immaterial. What matters is that the exposure is linear in both the level and the change of emissions.

Let M_{t+1} denote the equilibrium stochastic discount factor in the incomplete-markets economy with heterogeneous climate concern and unspanned damages, as in the baseline model. The conditional Euler equation implies that

$$0 = E_t[M_{t+1} R_{i,t+1}] - 1 = E_t[M_{t+1}](1 + R_{f,t}) + \mathrm{Cov}_t(M_{t+1}, R_{i,t+1}),$$

and a standard log-linear approximation yields

$$E_t[R_{i,t+1} - R_{f,t}] \;=\; \lambda_{\zeta,t}\, g_\zeta(Z_{i,t}) + \text{terms orthogonal to } \zeta_{t+1},$$

where $\lambda_{\zeta,t}$ is the equilibrium price of transition risk associated with the factor ζ_{t+1}. Since $g_\zeta(Z_{i,t})$ is linear in $(X_{i,t}, \Delta X_{i,t})$, we obtain the conditional cross-sectional relation

$$E_t[R_{i,t+1} - R_{f,t}] \;=\; \lambda_{\zeta,t}(\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}) + \text{other factor terms}. \tag{6.7}$$

Equation (6.7) shows that both emission levels and changes appear as natural priced characteristics once the size of the transition loss depends on how far and how fast a firm is drifting away from a compatible transition path. In the special case $\phi_2 = 0$, transition losses depend only on the level of emissions and the model collapses back to the level-only specification. When $\phi_2 > 0$, firms that are increasing their emissions are more exposed to an abrupt tightening shock, and a regression of returns on both $X_{i,t}$ and $\Delta X_{i,t}$ identifies the linear combination $\phi_1 X_{i,t} + \phi_2 \Delta X_{i,t}$ implied by the

transition-loss technology.

This construction modifies only the $g_\zeta(Z_{i,t})$ block of the model. The incomplete-markets structure, the heterogeneous climate concern across investors, and the decomposition of the transition price of risk into financial and non-pecuniary components are unchanged. The result is that emission growth becomes a natural priced characteristic alongside emission levels, providing an asset-pricing foundation for empirical specifications that use both levels and changes of emissions as proxies for exposure to carbon-transition risk.

7 Jevons-type Efficiency Effects

This section illustrates how endogenous energy efficiency can further weaken the case for any single emission-based statistic to be a unique sufficient summary of climate exposure in equilibrium returns. The goal is not to build a full production model, but to show that once firms choose an efficiency parameter that shifts energy demand, the *priced characteristics* in the linear factor structure become joint functions of efficiency and emissions. A Jevons-type rebound effect then implies that higher efficiency need not reduce energy use or emissions and can, for some firms, *raise* equilibrium risk premia by increasing their effective exposure to priced climate or policy shocks. Bolton et al. (2025) provide recent evidence of this phenomenon in the context of green patents and emission levels.

Efficiency choice and rebound. At date t, firm i uses energy $E_{i,t}$ and 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 the first-order condition

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

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 $X_{i,t}$ is itself an equilibrium object induced by the choice $E_{i,t}^*(a_{i,t})$, 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 the sign of its response to efficiency need not be negative in the rebound region.

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: $a_{i,t} \mapsto E_{i,t}^*(a_{i,t}) \mapsto X_{i,t}$, so the exposure coefficients in (7.1) generically depend on $(X_{i,t}, a_{i,t})$ jointly.

Under the same linear-factor pricing logic as in Proposition 4 and Section 4, expected excess returns admit the approximation

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

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 it immediate that the priced characteristics are the exposure functions $g_\zeta(\cdot)$ and $g_E(\cdot)$, which depend on efficiency and emissions jointly; there is no generic reason for premia to collapse to a function of $X_{i,t}/V_{i,t}$ (or to $X_{i,t}$ alone) once $a_{i,t}$ shifts energy use and hence shock sensitivity.

A simple example with reversed comparative statics

To see how a rebound can raise premia, consider a simple service-demand specification

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

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

$$E^*_{i,t}(a_{i,t}) = \frac{a_{i,t}p - P^E_t}{ca^2_{i,t}} \quad \text{for } a_{i,t}p > P^E_t,$$

so that

$$\frac{\partial E^*_{i,t}}{\partial a_{i,t}} = \frac{2P^E_t - a_{i,t}p}{ca^3_{i,t}}.$$

Hence, for $a_{i,t} \leq 2P^E_t/p$, 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 $\lambda^{(\zeta)}_t > 0$ (for example, when the financial brown-premium component dominates the taste-driven component). Then (7.2) reduces to

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

and therefore, in the rebound region,

$$\frac{\partial}{\partial a_{i,t}}\left(\mathbb{E}_t[R_{i,t+1}] - R_{f,t}\right) = \lambda^{(\zeta)}_t \eta\gamma \frac{\partial E^*_{i,t}}{\partial a_{i,t}} \geq 0,$$

with strict inequality for $a_{i,t} < 2P^E_t/p$. In words, higher efficiency can increase equilibrium risk premia by raising emissions and hence exposure to the priced climate factor.

Relation to the planner wedge and to intensity sufficiency. In the complete-markets planner benchmark, Jevons-type responses are internalized in the optimal paths for $(a^*_{i,t}, E^*_{i,t})$ and therefore affect the *levels* of optimal emissions and the SCC, but they do not alter the basic sustainable-discount-rate representation in which the deterministic implementation wedge is proportional to $X^*_{i,t}/V^*_{i,t}$. By contrast, in equilibrium asset pricing with heterogeneous agents and

incomplete spanning, Jevons operates by reshaping the exposure functions $g_\zeta(X_{i,t}, a_{i,t}, Z_{i,t})$ and $g_E(X_{i,t}, a_{i,t}, Z_{i,t})$ that enter (7.2). Because these exposures depend on efficiency and emissions jointly and need not be homogeneous in $(X_{i,t}, V_{i,t})$, the conditions required for emission intensity $X_{i,t}/V_{i,t}$ to be a unique sufficient statistic are generically violated. The Jevons mechanism thus reinforces the paper's central theme that outside the complete-markets planner benchmark, priced climate characteristics are model-dependent exposure functions, and efficiency improvements can be associated with higher required returns for some firms.

8 Welfare and Empirical Implications

8.1 Planner wedges vs equilibrium premia

Complete markets with optimal policy In a complete-markets setting with a representative agent and a full span of climate risk, there exists a social planner who chooses consumption and emissions paths to maximize discounted welfare. Let C_t denote aggregate consumption at date t, let $\bar{X}_t$ denote aggregate emissions, and let Γ_t denote aggregate damages. The planner solves

$$\max_{\{C_t, \bar{X}_t, \dots\}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t u\big(C_t - \Phi\Gamma_t\big)\right], \tag{8.1}$$

subject to the technological and resource constraints of the economy. The scalar $\Phi > 0$ is a normative weight on climate damages that may, but need not, coincide with any particular aggregation of private climate concerns. The first-order conditions of the planner's problem define a sequence of social costs of carbon $\{S_t\}_t$, where S_t is the shadow value, in units of the consumption good, of a marginal increase in emissions at date t:

$$S_t \equiv \frac{\partial \mathcal{W}/\partial \bar{X}_t}{\partial \mathcal{W}/\partial C_t}, \tag{8.2}$$

with $\mathcal{W}$ denoting the planner's objective in (8.1). Intuitively, S_t measures the marginal welfare loss from an additional unit of emissions, expressed in consumption-equivalent units.

A standard welfare result implies that the planner's allocation can be implemented as a competitive equilibrium with a Pigouvian tax τ_t per unit of emissions, set equal to the SCC: $\tau_t = S_t$ for all t. In the decentralized economy, a profit-maximizing firm that takes $\{\tau_t\}_t$ as given chooses

its emission path to satisfy

$$\text{marginal private benefit of emitting} = \text{marginal private cost} + \tau_t.$$

Setting $\tau_t = S_t$ makes the firm's emission first-order condition coincide with the planner's condition, so that the competitive allocation with the Pigouvian tax $\{\tau_t\}_t$ reproduces the planner's allocation.

In an equivalent representation based on sustainable discount rates, let $V^*_{i,t}$ denote the value of firm i at date t along the planner's allocation and let $X^*_{i,t}$ denote its emissions. For a given financial discount rate r_t, one can define firm-specific sustainable discount rates $r^x_{i,t}$ such that discounting untaxed cash flows at $r^x_{i,t}$ yields the same value as discounting tax-adjusted cash flows at r_t. To a first-order approximation, the wedge between the financial and sustainable discount rates takes the form

$$r^x_{i,t} \approx r_t + (S_t - \tau_t)\frac{X^*_{i,t}}{V^*_{i,t}}, \tag{8.3}$$

where the deviation $(S_t - \tau_t)$ captures the gap between the prevailing carbon price and the SCC. When $\tau_t = S_t$, the sustainable and financial discount rates coincide along the planner's allocation, and the first–best allocation is decentralized either via Pigouvian taxes or via sustainable discount rates that internalize the SCC through the emission intensity $X^*_{i,t}/V^*_{i,t}$. Relatedly, Chittaro et al. (2025) show that in a frictionless representative-agent economy a carbon tax can be replicated by a return schedule that increases linearly in firms' emission intensity, and use this equivalence to quantify the pollution premia required to mimic realistic tax levels.

8.2 Incomplete markets and partial internalization

Consider now the dynamic incomplete-markets economy described in Section 7, with households indexed by h, preferences

$$U^h = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t u^h(C_{h,t} - \phi_h \Gamma_t)\right], \tag{8.4}$$

and climate damages

$$\Gamma_t = \kappa \bar{X}_{t-1}\zeta_t, \qquad \bar{X}_{t-1} = \sum_i X_{i,t-1}, \tag{8.5}$$

where $\phi_h \geq 0$ captures the climate concern of household h. Let $(\{C_{h,t}\}_{h,t}, \{X_{i,t}\}_{i,t}, \{K_{i,t}\}_{i,t})$ denote a feasible allocation for this economy and let $\{P_{i,t}\}_{i,t}$ and $\{R_{f,t}\}_t$ be asset prices that support a competitive equilibrium with the given preferences and technology. Climate damages remain non-

tradable, so markets are incomplete: the consumption risk generated by Γ_t cannot be fully hedged via traded assets.

For any household h, the intertemporal first-order condition for optimal portfolio choice implies the existence of a household-specific stochastic discount factor (SDF)

$$m_{h,t+1} = \beta \frac{u_c^h(C_{h,t+1} - \phi_h \Gamma_{t+1})}{u_c^h(C_{h,t} - \phi_h \Gamma_t)}, \tag{8.6}$$

such that

$$P_{i,t} = \mathbb{E}_t\big[m_{h,t+1} D_{i,t+1}\big], \qquad \mathbb{E}_t[m_{h,t+1} R_{i,t+1}] = 1, \tag{8.7}$$

for all i and all h who trade asset i. Under standard regularity conditions, there exists an *aggregate* SDF m_{t+1} that prices all traded assets,

$$P_{i,t} = \mathbb{E}_t\big[m_{t+1} D_{i,t+1}\big], \qquad \mathbb{E}_t[m_{t+1} R_{i,t+1}] = 1, \tag{8.8}$$

and can be written as a wealth-weighted aggregation of the individual SDFs $\{m_{h,t+1}\}_h$. Because markets are incomplete and climate concern is heterogeneous, this aggregate SDF does not coincide with the intertemporal marginal rate of substitution of any social planner who assigns a single normative weight Φ to climate damages. In particular, even if Φ is chosen as some average of the $\{\phi_h\}_h$, the planner's SDF

$$m_{t+1}^{\mathrm{SP}} = \beta \frac{u_c(C_{t+1}^* - \Phi \Gamma_{t+1}^*)}{u_c(C_t^* - \Phi \Gamma_t^*)}$$

that supports the solution to (8.1) is, in general, distinct from m_{t+1}.

In the CARA-normal environment described in Section 7, the aggregate SDF m_{t+1} admits the exponential-affine representation

$$m_{t+1} = \exp\big(-A_t - B_t \zeta_{t+1}\big), \tag{8.9}$$

for suitable processes $\{A_t\}_t$ and $\{B_t\}_t$, where $B_t > 0$ reflects both the cross-section of risk aversion $\{\alpha_h\}_h$ and the aggregate climate concern $\sum_h \phi_h$. A second-order (log-linear) approximation to the Euler equation (8.8), together with the linear exposure of dividends to ζ_{t+1}, yields the dynamic pricing relation

$$\mathbb{E}_t\big[R_{i,t+1}\big] - R_{f,t} \approx \lambda_{X,t} X_{i,t}, \tag{8.10}$$

with

$$\lambda_{X,t} = \frac{\sigma_\zeta^2 \bar{X}_t \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} - \frac{\sigma_\zeta^2 \bar{X}_t \kappa \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \sum_{h=1}^{H} \phi_h. \tag{8.11}$$

The first term in (8.11) is the purely financial price of unhedgeable climate risk: it is increasing in aggregate emissions $\bar{X}_t$, in the volatility of the climate shock σ_ζ^2, and in the technological sensitivity η, and decreasing in aggregate risk tolerance $\sum_h (1/\alpha_h)$. The second term captures the compression of premia induced by investors who directly dislike damages: as aggregate climate concern $\sum_h \phi_h$ rises, climate-exposed assets become less attractive and the equilibrium price of climate risk is reduced. Crucially, the firm-specific exposure in (8.10) is the *level* of emissions $X_{i,t}$, not the ratio $X_{i,t}/V_{i,t}$, and the coefficient $\lambda_{X,t}$ is determined by the equilibrium distribution of $\{\alpha_h, \phi_h\}_h$ under incomplete risk sharing.

By contrast, in the complete-markets benchmark with a representative planner who assigns the normative weight Φ to damages, the planner's problem (8.1) delivers a sequence of SCCs $\{S_t\}_t$ and first-best optimality conditions that can be implemented via either Pigouvian taxes $\tau_t = S_t$ or intensity-based sustainable discount rates of the form (8.3). In that environment, the wedge between financial and sustainable discount rates is uniquely characterized by the product of the SCC and the emission intensity $X_{i,t}^*/V_{i,t}^*$, and the planner's SDF m_{t+1}^{SP} coincides with the unique SDF that prices all assets.

In the incomplete-markets equilibrium described by (8.9)-(8.11), the discount rates actually faced by firms are governed by the term $\lambda_{X,t} X_{i,t}$. The climate price of risk $\lambda_{X,t}$ is endogenous to the cross-section of $\{\alpha_h, \phi_h\}_h$ and aggregate emissions $\bar{X}_t$, and there is, in general, no simple mapping from $\lambda_{X,t}$ to the planner's SCC S_t. Absent a policy instrument, the privately optimal paths $\{X_{i,t}, K_{i,t}\}_{i,t}$ that arise from (8.10) do not satisfy the planner's first-order conditions for emissions and capital, except in degenerate cases in which the heterogeneity in ϕ_h is eliminated and markets are effectively complete.

It follows that the competitive equilibrium with heterogeneous investors and incomplete markets is a second-best allocation in which climate damages are only partially internalized. Part of the externality is reflected in asset prices through the true greenium component of $\lambda_{X,t}$, driven by $\sum_h \phi_h$, but the equilibrium SDF remains distinct from the planner's SDF and the implied discount-rate wedges differ from those associated with the SCC. The complete-markets allocation supported by Pigouvian taxes or by $X_{i,t}/V_{i,t}$-based sustainable discount rates is therefore generically welfare-superior to the incomplete-markets equilibrium, even though both frameworks feature well-defined

prices of climate risk and well-defined emission-based risk characteristics.

8.3 Welfare losses from market incompleteness

The analysis above implies that, once climate damages are unspanned and investors differ in climate concern, equilibrium allocations are generically second best relative to the complete-markets planner benchmark. It is therefore useful to make precise what is meant by a welfare loss from market incompleteness in a way that is distinct from, and not directly inferred from, observed carbon premia.

Let $W(\cdot)$ denote the planner welfare functional from Section 3, so that for any feasible aggregate allocation (C_0, C_1, X_1),

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

Let $(C_0^{FB}, C_1^{FB}, X_1^{FB})$ denote the first-best planner allocation and define

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

Let $(C_0^{IM}, C_1^{IM}, X_1^{IM})$ denote the aggregate allocation induced by the incomplete-markets equilibrium (with the same technology, damage function, and shock distribution), and define the planner evaluation of this second-best allocation as

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

The welfare loss from market incompleteness can then be expressed in consumption-equivalent units. Specifically, define the date-1 consumption-equivalent loss $\Delta c_1 \geq 0$ as the unique scalar (on the relevant domain) that solves

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

We then report the welfare loss either in levels, Δc_1, or as a fraction of first-best date-1 consumption,

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

This definition makes it clear that the welfare loss is an allocation object, the consumption reduction

that makes the planner indifferent between the first-best allocation and the incomplete-markets equilibrium, and it need not be monotone in the magnitude of equilibrium carbon premia. Carbon premia summarize second-best pricing of climate exposures in equilibrium, whereas ΔW is pinned down by the gap between allocations under complete versus incomplete risk sharing, evaluated under a common welfare criterion.

In the present framework, ΔW captures two distinct channels. First, incomplete markets prevent the economy from fully insuring aggregate climate risk, so consumption allocations differ from the planner optimum even absent preference heterogeneity. Second, heterogeneous climate concern implies that the equilibrium stochastic discount factor reflects a weighted aggregation of private marginal rates of substitution rather than the planner's marginal valuation of damages. As a result, the equilibrium price of climate risk generically deviates from the social cost of carbon, leading to misaligned investment incentives and an inefficient allocation of capital across firms with different emission profiles.

Importantly, ΔW is not directly inferred from observed carbon premia. A large or small greenium does not, by itself, imply a correspondingly large or small welfare loss, because equilibrium premia reflect second-best pricing under incomplete markets rather than planner wedges. This distinction clarifies why carbon premia should not be interpreted as sufficient statistics for the welfare impact of climate risk or climate policy. Instead, welfare comparisons must be conducted relative to an explicit planner benchmark that internalizes damages and spans aggregate risk, even when equilibrium pricing relations take a similar linear factor form.

8.4 Empirical implications and discussion

The incomplete-markets structure developed above has testable implications for the cross-section of expected returns and for the interpretation of existing evidence on brown premium and greenium. This section sketches the main empirical predictions and contrasts them with those implied by the complete-markets sustainable-discount-rate benchmark.

More generally, the CARA-normal incomplete-markets environment with generic dividend exposure

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

implies the conditional linear pricing relation

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \sum_{k=1}^{K} \lambda_t^{(k)} g_k(Z_{i,t}), \tag{8.12}$$

where $Z_{i,t}$ collects firm characteristics and state variables, and $\lambda_t^{(k)}$ are the prices of the K climate and policy risk factors. Equation (8.12) maps into the generic cross-sectional regression

$$R_{i,t+1} - R_{f,t} = \alpha_t + \sum_{k=1}^{K} \beta_t^{(k)} g_k(Z_{i,t}) + \varepsilon_{i,t+1}, \tag{8.13}$$

with $\beta_t^{(k)}$ estimating $\lambda_t^{(k)}$ up to sampling error and potential misspecification. The choice of the functions g_k constitutes the empirical analogue of specifying how climate and policy shocks enter technology and damages. Emission levels $X_{i,t}$, emission intensities $X_{i,t}/V_{i,t}$, changes $\Delta X_{i,t}$, and interacted terms (for example, $X_{i,t}$ multiplied by industry indicators) all arise as special cases of $g_k(Z_{i,t})$.

By contrast, in the complete-markets, representative-agent benchmark underlying the sustainable-discount-rate approach, the planner's problem yields a social cost of carbon S_t and sustainable discount rates

$$r_{i,t}^{x} = r_t + (S_t - \tau_t)\frac{X_{i,t}}{V_{i,t}},$$

where τ_t is the carbon tax and $V_{i,t}$ is firm value. In that environment, and abstracting from traditional risk factors, the relevant cross-sectional wedge in expected returns is proportional to the *intensity* $X_{i,t}/V_{i,t}$ rather than to the level $X_{i,t}$ itself. A natural empirical implementation of the sustainable-discount-rate logic therefore takes the form

$$R_{i,t+1} - R_{f,t} = \alpha_t + \gamma_t \frac{X_{i,t}}{V_{i,t}} + \varepsilon_{i,t+1}, \tag{8.14}$$

or, more generally, regressions in which $X_{i,t}/V_{i,t}$ is included as a characteristic or factor and the tax–adjusted component $(S_t - \tau_t)$ is interpreted through the magnitude of γ_t. In a strict interpretation of this benchmark, once standard risk factors and $X_{i,t}/V_{i,t}$ are controlled for, emission levels $X_{i,t}$ should not carry independent explanatory power for expected excess returns.

The incomplete-markets model thus yields a different set of empirical predictions. In one tractable specification, expected returns load on emission levels as in (8.17); in the general specification (8.13), the priced emission-based objects are the functions $g_k(Z_{i,t})$ that govern dividend

exposure to climate and policy shocks, which need not coincide with $X_{i,t}/V_{i,t}$. This implies that, in empirical work, it is natural to estimate specifications that include both $X_{i,t}$ and $X_{i,t}/V_{i,t}$ (and potentially other emission-based transformations) and to interpret the significance and signs of their coefficients through the lens of incomplete risk sharing and heterogeneous climate concern. A finding that emission levels carry a robust premium controlling for intensity is particularly natural in the incomplete-markets framework and difficult to reconcile with a strict sustainable-discount-rate interpretation in which $X_{i,t}/V_{i,t}$ is the uniquely relevant object.

As a special case, in the static baseline model with heterogeneous CARA investors, unspanned climate damages, and dividend exposure

$$D_i = \mu_i + \eta X_i \zeta,$$

Proposition 2 implies

$$\mathbb{E}[R_i] - R_f = \lambda \eta X_i, \tag{8.15}$$

with

$$\lambda = \frac{\sigma_\zeta^2 \bar{X}}{\sum_h \frac{1}{\alpha_h}} \Big(\eta - \kappa \sum_h \varphi_h\Big),$$

and the decomposition $\lambda = \lambda^{\text{fin}} + \lambda^{\text{green}}$ given earlier. In a dynamic version, the corresponding conditional relation is

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_{X,t} X_{i,t}, \tag{8.16}$$

where $\lambda_{X,t}$ is an endogenous function of the cross-section of (α_h, φ_h) and of aggregate emissions $\bar{X}_t$. Equations (8.15)-(8.16) lead directly to the panel regression specification

$$R_{i,t+1} - R_{f,t} = \alpha_t + \beta_{X,t} X_{i,t} + \varepsilon_{i,t+1}, \tag{8.17}$$

possibly augmented with standard risk factors and controls, where under the maintained assumptions the conditional expectation of the slope satisfies

$$\mathbb{E}_t[\beta_{X,t}] \approx \lambda_{X,t}.$$

In particular, in the absence of climate concern ($\varphi_h = 0$ for all h) the model predicts a positive emission-level slope $\beta_{X,t} > 0$ reflecting a purely financial brown premium for bearing unhedgeable

climate risk. When some investors are green ($\sum_h \varphi_h > 0$), the true greenium component λ^{green} partially offsets this financial premium and reduces $\beta_{X,t}$; for sufficiently strong aggregate climate concern, $\beta_{X,t}$ can even become negative.

8.5 Empirical evidence on sufficiency

To illustrate how the theoretical framework maps into empirical specifications, I briefly consider a simple cross-sectional application using firm-level emissions and stock returns. The goal of this subsection is not to provide a definitive empirical contribution, but rather to show how the sufficiency logic developed above can be taken to the data and how level-, intensity-, and change-based measures of emissions can be interpreted through the lens of my model.

I collect panel data on stock returns $R_{i,t+1}$, the risk-free rate $R_{f,t}$, firm-level emissions $X_{i,t}$, firm value $V_{i,t}$, and standard firm characteristics $W_{i,t}$ (such as size, profitability, investment, and momentum) over periods $t = 1, \ldots, T$ and firms $i = 1, \ldots, N_t$. Following Pedersen (2025), I define emission intensity as $X_{i,t}/V_{i,t}$, where V is proxied for by the book value of equity,[1] and the change in emissions as $\Delta X_{i,t} = X_{i,t} - X_{i,t-1}$. Guided by the conditional pricing relation

$$E_t[R_{i,t+1} - R_{f,t}] = \lambda_{X,t} X_{i,t}$$

in the incomplete-markets benchmark and the sufficiency analysis in the factor framework, we estimate cross-sectional regressions of the form

$$R_{i,t+1} - R_{f,t} = \alpha_t + \beta_{X,t} X_{i,t} + \gamma_t^\top W_{i,t} + \varepsilon_{i,t+1}, \tag{8.18}$$

$$R_{i,t+1} - R_{f,t} = \alpha_t + \beta_{I,t} \frac{X_{i,t}}{V_{i,t}} + \gamma_t^\top W_{i,t} + \varepsilon_{i,t+1}, \tag{8.19}$$

and

$$R_{i,t+1} - R_{f,t} = \alpha_t + \beta_{X,t} X_{i,t} + \beta_{I,t} \frac{X_{i,t}}{V_{i,t}} + \gamma_t^\top W_{i,t} + \varepsilon_{i,t+1}. \tag{8.20}$$

In the empirical implementation, I use a large sample of international firms over the period 2005-2022 and estimate pooled panel regressions of one-month-ahead and twelve-month-ahead excess returns on log Scope 1 emissions, Scope 1 emission intensity, and controls, including industry×time and country fixed effects.[2] Standard errors are double clustered by industry and time, and coeffi-

[1] This specification mimics closely Pedersen (2025) but the results are also robust to other scaling variables, including book value of assets, market capitalization, and sales.

[2] My regression specification largely follows that in Bolton and Kacperczyk (2023). In the literature on carbon

cients are reported along with their robust standard errors.

The interpretation of these regressions is directly tied to the theory. In the complete-markets planner benchmark with social-cost-of-carbon logic, emission intensity $X_{i,t}/V_{i,t}$ is the sufficient statistic for the sustainable-discount-rate wedge. In that environment, conditional on intensity, emission levels $X_{i,t}$ should not carry an additional premium. Once the sustainable discount rate is correctly adjusted for intensity, two firms with the same $X_{i,t}/V_{i,t}$ should not systematically differ in expected returns on account of their absolute emissions. In contrast, in the incomplete-markets economy with heterogeneous climate concern and unspanned damages, the equilibrium carbon premium is linear in emission levels $X_{i,t}$, and there is no theoretical reason for intensity to span the relevant exposure. In that environment, the slope of $X_{i,t}$ in (8.18) is an estimate of the conditional climate price of risk $\lambda_{X,t}$, and the intensity term in (8.19) is, at best, an imperfect proxy. The joint specification (8.20) provides a simple sufficiency test: if intensity were the empirically relevant object in the sense of Definition 4.1, the coefficient of $X_{i,t}$ should be close to zero once $X_{i,t}/V_{i,t}$ is included. Conversely, if levels are the relevant exposure, one would expect $X_{i,t}$ to remain priced while $X_{i,t}/V_{i,t}$ becomes redundant.

The results in Table 1 present the estimates of such specification. When entered on its own, emission intensity is positively related to one-month-ahead returns, whereas log Scope 1 emissions are strongly positively related to both one-month and twelve-month ahead returns. In the joint specifications, however, the coefficient of log emissions remains positive and highly statistically significant at both horizons, while the coefficient of intensity becomes economically and statistically indistinguishable from zero. The explanatory power of the regressions is essentially unchanged across columns. Taken together, these results suggest that absolute emissions are the more robust priced characteristic in the cross-section, and that intensity does not span the relevant carbon exposure once levels are controlled for, a configuration that is difficult to reconcile with a strict sustainable-discount-rate interpretation but fits naturally within the incomplete-markets framework.

To connect more directly to the transition-risk channel in Section 6, we also estimate a specification that allows for the possibility that both levels and changes of emissions enter the exposure

emissions there is an issue of properly specifying investors' information sets, and the risk of either under- or overconditioning. In the absence of a clear guidance for investors' information set I consider information lag of one month and one year, which are likely at the realistic extremes for most companies. Thus, the true information set for each firm-emission is somewhere in between the two numbers but with the lack of reliable data difficult to pin down.

function $g_\zeta(Z_{i,t})$. Specifically, we consider

$$R_{i,t+1} - R_{f,t} = \alpha_t + \beta_{X,t} X_{i,t} + \beta_{\Delta X,t} \Delta X_{i,t} + \gamma_t^\top W_{i,t} + \varepsilon_{i,t+1}, \tag{8.21}$$

as reported in Table 2. In the transition-loss specification $\Gamma_{i,t} = \varphi_1 X_{i,t} + \varphi_2 \Delta X_{i,t}$, the dividend exposure to an unexpected tightening shock is proportional to $\varphi_1 X_{i,t} + \varphi_2 \Delta X_{i,t}$, and the corresponding expected-return relation is

$$E_t[R_{i,t+1} - R_{f,t}] = \lambda_{\zeta,t}(\varphi_1 X_{i,t} + \varphi_2 \Delta X_{i,t}) + \text{other factor terms.}$$

Regression (8.21) is therefore a direct empirical implementation of this exposure structure.

The estimates in Table 2 indicate that both levels and changes of emissions are rewarded over short horizons. For one-month-ahead returns, log Scope 1 emissions are positively and significantly related to future returns when included on their own, and the coefficient remains positive and significant when emissions growth is added. Emissions growth itself enters with a positive and statistically significant coefficient, so that firms with higher current emissions and firms that are increasing their emissions more rapidly both earn higher short-horizon expected returns, conditional on standard controls. At the twelve-month horizon, by contrast, the growth term loses statistical significance and the premium on emission levels becomes smaller, though it remains positive and significant. This pattern is consistent with a transition-loss technology in which both levels and recent changes in emissions matter for near-term exposure to tightening shocks, but in which the effect of changes is more transitory than the effect of the level.

The robust significance of both emission levels and growth rates in Table 2 provides a direct empirical rejection of the intensity-sufficiency hypothesis. As derived in the knife-edge proof of Section 5, if intensity were the unique sufficient statistic for climate risk, neither the absolute scale of emissions nor their growth rates should carry incremental explanatory power in a well-specified asset-pricing relation. The fact that the market prices these characteristics independently suggests that the restrictive homogeneity conditions required for intensity to summarize exposure are not met. Instead, the results support a more general exposure structure where absolute fundamentals and their dynamics are the primary drivers of the carbon premium.

Overall, these simple regressions are only illustrative, and the coefficients should not be over-interpreted given the usual concerns about measurement error in emissions data, sample composition, and omitted risk factors. Nevertheless, they show how the theoretical distinctions developed

above can be brought to the data. In our sample, emission levels remain robustly priced after controlling for intensity and for emission growth, while intensity becomes redundant once levels are included and growth matters primarily at short horizons. This configuration is exactly what one would expect in an incomplete-markets economy with transition losses tied to emissions, and it is difficult to reconcile with a strict sustainable-discount-rate view in which intensity alone is the unique sufficient statistic for carbon exposure.

9 Conclusion

This paper provides a unified framework to reconcile the debate over whether carbon premia load on emission levels or intensities. By mapping firm fundamentals into an equilibrium exposure function, I show that the 'intensity-only' rule common in the social-planner literature is a restrictive, knife-edge case. In the presence of market incompleteness and unspanned damages, absolute emission levels emerge as the natural priced characteristic. This reconciliation suggests that the conflict in the empirical literature is not a matter of right or wrong metrics, but a reflection of the gap between normative policy benchmarks and positive market pricing.

A key contribution of the dynamic analysis is the identification of a technological sign-reversal. Through a Jevons-type rebound effect, I demonstrate that firm-level efficiency improvements can endogenously increase absolute emissions and, consequently, equilibrium risk premia. This result cautions against the use of intensity as a singular metric for greenness, as it may mask an underlying increase in a firm's fundamental risk exposure.

Beyond the choice of carbon metrics, this framework provides a micro-foundation for the green and brown classifications central to the climate finance literature (e.g., Pastor et al. (2021)). By endogenizing these types as equilibrium exposure functions rather than exogenous labels, I show that a firm's climate risk profile is a dynamic choice shaped by its technological environment. Specifically, the Jevons-type rebound effect demonstrates a technological sign-reversal. Improvements in energy efficiency intended to green a firm's profile can endogenously increase its absolute emissions and equilibrium risk premium. This implies that effective risk assessment must move beyond static ESG labels toward a functional understanding of how technology and fundamentals co-evolve to shape risk exposure.

Finally, the analysis offers a protocol for empirical research. Level-based specifications are not mere alternatives to intensity but are theoretically grounded signatures of incomplete markets.

Future work should focus on disentangling the brown-premium and greenium components of the climate price of risk and integrating these asset-pricing frictions into larger integrated assessment models (IAMs).

References

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.

Bolton, P., M. Kacperczyk, and M. Wiedemann (2025). The co2 question: Technical progress and the climate crisis. *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*.

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.

Tables

Table 1: Predicting stock returns with emissions levels and intensity

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Stock returns					
	1 month ahead			12 months ahead		
XV (Scope1)*100	0.019**		-0.000	0.006		0.000
	(0.009)		(0.011)	(0.010)		(0.011)
Log(Scope1 Emissions)		0.123***	0.123***		0.039***	0.039***
		(0.019)	(0.024)		(0.009)	(0.012)
Log(Size)	-0.200***	-0.269***	-0.269***	-0.139***	-0.162***	-0.162***
	(0.049)	(0.052)	(0.052)	(0.037)	(0.037)	(0.037)
Log(M/B)	-0.250***	-0.193**	-0.193**	-0.159*	-0.141*	-0.141*
	(0.072)	(0.068)	(0.072)	(0.076)	(0.072)	(0.077)
Book leverage	-0.060	-0.189	-0.189	-0.117	-0.156	-0.156
	(0.240)	(0.241)	(0.242)	(0.174)	(0.174)	(0.176)
Momentum	0.472***	0.458**	0.458**	-0.394*	-0.397*	-0.397*
	(0.160)	(0.161)	(0.161)	(0.190)	(0.190)	(0.190)
Investment/Assets	-0.408	-0.241	-0.241	0.099	0.153	0.153
	(0.865)	(0.877)	(0.891)	(0.738)	(0.752)	(0.757)
Log(PPE)	0.078**	0.041	0.041	0.056**	0.045**	0.045**
	(0.028)	(0.026)	(0.027)	(0.020)	(0.020)	(0.021)
Volatility	1.995	1.995	1.995	0.408	0.403	0.403
	(3.151)	(3.147)	(3.150)	(1.525)	(1.520)	(1.524)
ROE	0.842***	0.770**	0.770**	0.548	0.527	0.527
	(0.259)	(0.268)	(0.270)	(0.313)	(0.315)	(0.317)
Age	0.881	0.626	0.626	-0.442	-0.513	-0.513
	(0.935)	(0.959)	(0.972)	(0.940)	(0.971)	(0.990)
Constant	2.055***	1.738***	1.738***	1.791***	1.689***	1.689***
	(0.445)	(0.469)	(0.472)	(0.292)	(0.296)	(0.299)
Industry×time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,006,908	1,006,023	1,006,023	862,872	862,251	862,251
R-squared	0.231	0.231	0.231	0.247	0.247	0.247

Robust standard errors in parentheses

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

Table 2: 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

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

A Online Appendix: 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-level 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

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

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 and captures private costs of emissions (abatement, regulatory compliance, etc.).

Emissions of firm i at date t are

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

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

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

where $I_{i,t}$ is investment at date t. We assume that $(K_{i,0})_{i=1}^N$ are given and that all choice variables are nonnegative.

Climate risk and dividends Let $\{\zeta_t\}_{t\geq 1}$ be an i.i.d. climate or policy shock process with

$$\mathbb{E}[\zeta_{t+1}|\mathcal{F}_t] = 0, \qquad \mathrm{Var}[\zeta_{t+1}|\mathcal{F}_t] = \sigma_\zeta^2 > 0,$$

and assume joint conditional normality of all shocks introduced below. Climate damages at $t + 1$ are specified as

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

so that damages scale with aggregate emissions at t and the realization of ζ_{t+1}.

Firm i's dividend at date $t+1$ is given by

$$D_{i,t+1} = Y_{i,t+1} - I_{i,t+1} + \eta X_{i,t}\zeta_{t+1} + \tilde{\varepsilon}_{i,t+1}, \tag{A.5}$$

where $\eta > 0$ is a technological parameter capturing the sensitivity of dividends to the climate shock, and $\tilde{\varepsilon}_{i,t+1}$ is an idiosyncratic shock independent of ζ_{t+1} conditional on $\mathcal{F}_t$. The specification (A.5) preserves the static structure in which climate exposure is proportional to the *level* of emissions $X_{i,t}$.

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

$$U^h = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t u^h(C_{h,t} - \phi_h \Gamma_t)\right], \qquad \beta \in (0,1), \tag{A.6}$$

where $C_{h,t}$ is consumption at date t, $u^h(\cdot)$ is a CARA utility index,

$$u^h(x) = -\exp(-\alpha_h x), \qquad \alpha_h > 0,$$

and $\phi_h \geq 0$ measures the intensity of climate concern (disutility from damages) for household h. Households are heterogeneous in (α_h, ϕ_h).

Households trade a risk-free asset with gross return $R_{f,t}$ and the N equity claims $\{D_{i,t+1}\}_{i=1}^N$. Damages $\{\Gamma_t\}$ are *not* traded. Markets are therefore incomplete: the consumption risk induced by the non-pecuniary damage term cannot be fully hedged using traded assets.

Under CARA preferences and conditional joint normality of all relevant shocks, dynamic portfolio choice admits a mean-variance representation. Let

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

denote the climate-adjusted consumption of household h. Standard arguments imply that, given the independence of ζ_{t+1} across time and the absence of intertemporal hedging motives under CARA-normality, the household's problem at each date t can be written as choosing current consumption and portfolios to maximize

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

subject to the dynamic budget constraint and given state variables at t. The objective (A.7) is the exact equivalent of (A.6) in the CARA-normal setting and reduces the dynamic problem to a sequence of static mean-variance problems, one for each t.

Aggregate stochastic discount factor and its climate loading Let $\theta_{i,t}^h$ denote the number of shares of firm i held by household h at time t, and let b_t^h denote the holdings of the risk-free asset. The portfolio return of household h from t to $t+1$ is

$$R_{p,t+1}^h = \frac{1}{W_t^h}\left(b_t^h R_{f,t} + \sum_{i=1}^N \theta_{i,t}^h D_{i,t+1}\right),$$

where W_t^h is beginning-of-period wealth. Using (A.4)-(A.5) and market clearing $\sum_h \theta_{i,t}^h = 1$, we can write the innovation to climate-adjusted consumption of household h as

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

where $\tilde{\varepsilon}_{h,t+1}$ collects idiosyncratic shocks orthogonal to ζ_{t+1}.

From (A.7)-(A.8), the first-order condition for $\theta_{i,t}^h$ implies, as in the static case, that

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

Since the left-hand side of (A.9) does not depend on h, equilibrium requires that

$$\frac{\mathbb{E}_t[R_{i,t+1}] - R_{f,t}}{\eta X_{i,t}} = \alpha_h \sigma_\zeta^2 \delta_{h,t}$$

is equal to a common scalar λ_t for all h with interior holdings. Aggregating as in the static case by summing the definition of $\delta_{h,t}$ across h and using market clearing, we obtain the dynamic analogue of Proposition 2:

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

To link this mean-variance representation to an SDF, suppose that the aggregate SDF between t and $t+1$ admits the exponential–affine form

$$m_{t+1} = M_t \exp(-B_t \zeta_{t+1}), \qquad B_t > 0, \tag{A.11}$$

for some $\mathcal{F}_t$-measurable scaling factor M_t, normalized such that $\mathbb{E}_t[m_{t+1}] = 1/(1 + R_{f,t})$. The representation (A.11) is exact under CARA-normal preferences when all relevant shocks are linear in ζ_{t+1}; more generally, it is a second-order approximation under the joint conditional normality assumption.

For an asset with payoff

$$D_{i,t+1} = \bar{D}_{i,t+1} + \eta X_{i,t}\zeta_{t+1},$$

where $\bar{D}_{i,t+1}$ is conditionally independent of ζ_{t+1}, the standard log-linear approximation to the Euler equation

$$P_{i,t} = \mathbb{E}_t[m_{t+1}D_{i,t+1}]$$

implies the approximate risk-premium formula

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx B_t\sigma_\zeta^2\eta X_{i,t}. \tag{A.12}$$

Comparing (A.12) with the mean-variance relation

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} = \lambda_t\eta X_{i,t}$$

and using (A.10) identifies the SDF loading on the climate shock as

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

Thus the SDF loading B_t is explicitly determined by the cross-section of risk aversion and climate concern (α_h, ϕ_h), as well as by aggregate emissions $\bar{X}_t$ and the technological parameters η and κ.

The derivation of (A.12)-(A.13) relies on the following conditions: (i) CARA utility, (ii) joint conditional normality of ζ_{t+1} and dividend shocks, (iii) linear exposure of dividends and damages to ζ_{t+1} as in (A.4)-(A.5), and (iv) a second-order (log-linear) expansion of the Euler equation around small conditional variance of ζ_{t+1}. Under these standard CARA-normal linearization arguments, the dynamic pricing relation inherits the same emission-level structure as the static model.

Dynamic asset pricing relation Combining (A.12) and (A.13), we obtain the dynamic counterpart to the static emission-level factor:

$$\mathbb{E}_t[R_{i,t+1}] - R_{f,t} \approx \lambda_{X,t} X_{i,t}, \qquad \lambda_{X,t} \equiv \frac{\sigma_\zeta^2 \bar{X}_t \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} - \frac{\sigma_\zeta^2 \bar{X}_t \kappa \eta}{\sum_{h=1}^{H} \frac{1}{\alpha_h}} \sum_{h=1}^{H} \phi_h. \tag{A.14}$$

Expression (A.14) is exactly the time-t analogue of the static formula in Proposition 2, with time variation in $\bar{X}_t$ and in the distribution of (α_h, ϕ_h) inducing time variation in the climate price of risk $\lambda_{X,t}$. The exposure is the level of emissions $X_{i,t}$, not intensity $X_{i,t}/V_{i,t}$.

Firms' dynamic investment and emission choices Let $P_{i,t}$ denote the ex-dividend price of firm i at time t. Shareholders value the stream of dividends using the aggregate SDF:

$$P_{i,t} = \mathbb{E}_t \left[\sum_{s=t+1}^{\infty} \left(\prod_{u=t+1}^{s} m_u \right) D_{i,s} \right]. \tag{A.15}$$

Firm i chooses $\{K_{i,t+1}, q_{i,t}\}_{t \geq 0}$ to maximize $P_{i,0}$, taking as given prices and the SDF process $\{m_{t+1}\}$.

The Bellman equation for firm i is

$$P_{i,t} = \max_{K_{i,t+1}, q_{i,t}} \left\{ \mathbb{E}_t[m_{t+1} D_{i,t+1}] + \mathbb{E}_t[m_{t+1} P_{i,t+1}] \right\},$$

subject to (A.1)-(A.3) and (A.5). The first-order condition with respect to $q_{i,t}$ is

$$\mathbb{E}_t \left[m_{t+1} \frac{\partial D_{i,t+1}}{\partial q_{i,t}} \right] + \mathbb{E}_t \left[m_{t+1} \frac{\partial P_{i,t+1}}{\partial q_{i,t}} \right] = 0. \tag{A.16}$$

Under the maintained assumption that $q_{i,t}$ affects dividends only through $Y_{i,t+1}$ and $X_{i,t}$, and does not directly affect continuation values beyond $t+1$, the second term in (A.16) vanishes and we obtain

$$\mathbb{E}_t \left[m_{t+1} \left(\frac{\partial F_i(K_{i,t}, q_{i,t})}{\partial q_{i,t}} - \Psi_i'(q_{i,t}) + \eta \chi_i \zeta_{t+1} \right) \right] = 0. \tag{A.17}$$

The first two terms inside the expectation in (A.17) are the marginal private revenue from $q_{i,t}$; the last term is the marginal contribution of $q_{i,t}$ to climate-related risk, through its effect on emissions $X_{i,t} = \chi_i q_{i,t}$ and hence on the exposure $\eta X_{i,t} \zeta_{t+1}$ in (A.5).

Using the exponential-affine form (A.11) and the normality of ζ_{t+1}, we compute

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

so that (A.17) can be written as

$$\mathbb{E}_t[m_{t+1}]\left(\frac{\partial F_i(K_{i,t},q_{i,t})}{\partial q_{i,t}}-\Psi_i'(q_{i,t})-\eta\chi_i B_t\sigma_\zeta^2\right)=0. \tag{A.18}$$

Dividing by $\mathbb{E}_t[m_{t+1}]$ and using (A.13) yields the condition

$$\frac{\partial F_i(K_{i,t},q_{i,t})}{\partial q_{i,t}}-\Psi_i'(q_{i,t})=\eta\chi_i B_t\sigma_\zeta^2=\eta\chi_i\lambda_t, \tag{A.19}$$

which shows that the marginal private return to emissions is equated to a climate-risk premium term proportional to λ_t. The right-hand side depends on aggregate emissions $\bar{X}_t$ and the cross-section of (α_h,ϕ_h) through (A.10), and involves the emission *level* via the coefficient χ_i. In a complete-markets benchmark with a representative planner who internalizes damages, the corresponding condition would involve a Pigouvian tax or sustainable discount-rate wedge proportional to the planner's SCC and to emission intensity $X_{i,t}/V_{i,t}$. In the incomplete-markets environment above, the actual discount-rate wedge faced by the firm is instead governed by $\lambda_t X_{i,t}$, as in (A.14), and generally does not coincide with the planner's $X_{i,t}/V_{i,t}$-based wedge, even when the same damage function is used in the planner's problem.

Under the CARA-normal linearization assumptions detailed above, the infinite-horizon production economy thus preserves the key asset-pricing result of the static model: expected excess returns load on emission levels, with a climate price of risk that is explicitly determined by the cross-section of risk aversion and climate concern and by aggregate emissions.

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 wealth-weighted risk tolerance rather than by the sum of inverse CARA coefficients.

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 $h = 1, \ldots, H$. 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 P_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

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

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:

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

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 $\bar{Y}_h$ denote a deterministic reference level for Y_h; for the local approximation we treat $\bar{Y}_h$ as fixed and 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 using $\mathbb{E}[\zeta] = 0$ implies

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

Since $\mathrm{Var}(Y_h) = B_h^2 \sigma_\zeta^2$, 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

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

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 $u_h'(\bar{Y}_h) > 0$ (which does not affect the argmax over portfolios) yields the equivalent objective

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

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

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

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

Using (B.1), (B.2), and (B.4), we have

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

Denote

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

The mean-variance objective (B.11) becomes

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

The first-order condition with respect to θ_i^h is

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

Equivalently,

$$\mathbb{E}[R_i] - R_f = \tilde{\alpha}_h \sigma_\zeta^2 \delta_h \eta X_i, \tag{B.13}$$

for all i and all h with interior holdings, where $\mathbb{E}[R_i] = \mu_i / P_i$ as before.

Since the left-hand side of (B.13) does not depend on h, equilibrium requires that

$$\frac{\mathbb{E}[R_i] - R_f}{\eta X_i} = \tilde{\alpha}_h \sigma_\zeta^2 \delta_h$$

be equal to a common scalar λ for all h and all i with interior holdings. That is,

$$\mathbb{E}[R_i] - R_f = \lambda \eta X_i, \qquad \delta_h = \frac{\lambda}{\tilde{\alpha}_h \sigma_\zeta^2}. \tag{B.14}$$

To solve for λ, multiply the definition of δ_h in (B.12) by η and sum over h:

$$\sum_{h=1}^{H} \sum_{j=1}^{N} \theta_j^h \eta^2 X_j - \kappa \bar{X} \eta \sum_{h=1}^{H} \phi_h = \eta \sum_{h=1}^{H} \delta_h. \tag{B.15}$$

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.15), 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 5 (Price of climate risk under CRRA)**.** *Under the assumptions of this appendix, in*

a competitive equilibrium with CRRA preferences and incomplete markets, the equilibrium price of climate risk is

$$\lambda = \frac{\sigma_\zeta^2 \bar{X}}{\sum_{h=1}^{H} \frac{1}{\tilde{\alpha}_h}} \Big(\eta - \kappa \sum_{h=1}^{H} \phi_h\Big), \qquad \tilde{\alpha}_h = \frac{\gamma_h}{\bar{Y}_h}. \tag{B.16}$$

Accordingly, expected excess returns satisfy

$$\mathbb{E}[R_i] - R_f = \lambda_X X_i, \qquad \lambda_X \equiv \lambda\eta, \tag{B.17}$$

so that emission levels X_i again act as the relevant exposure in the one-factor structure.

The denominator in (B.16) 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

Proposition 5 shows that, in a CRRA-incomplete-markets setting, the local linear pricing relation

$$\mathbb{E}[R_i] - R_f = \lambda_X X_i$$

continues to hold, with a price of climate risk λ given by (B.16). The structure of the result is identical to the CARA-normal benchmark: the exposure is 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.

This derivation relies on a local approximation around reference levels $\{\bar{Y}_h\}_h$ at which the CRRA utility can be represented by a quadratic mean-variance objective with effective absolute risk aversion coefficients $\tilde{\alpha}_h = \gamma_h/\bar{Y}_h$. Under this 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 an X-factor representation in the cross-section of returns. 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 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 remain linear in firm-specific exposure coefficients, and the relevant priced characteristics are generic functions $g_k(Z_{i,t})$ of firm characteristics rather than a uniquely determined emission-intensity statistic $X_{i,t}/V_{i,t}$.

C.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

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

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,

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

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

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

where h_ζ and $h_{\zeta\zeta}$ denote the first and second partial derivatives with respect to ζ. Substituting (C.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

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

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

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

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

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

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

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

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 (C.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 function $g(Z_i)$ to a vector $(g_1(Z_i), g_2(Z_i))$, 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 (C.7) implies that expected excess returns are proportional to X_i with a combined price of risk $\lambda^{(1)}\eta_1 + \lambda^{(2)}\eta_2$, and the qualitative conclusion that emission levels act as the relevant exposure remains unchanged. More generally, the quadratic damage term induces an additional priced characteristic $g_2(Z_i)$ 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.

C.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 5, suppose that firm i's dividend can be written as

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

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

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

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

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

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 (C.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

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

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 $H_{k\ell}(\bar{X}_t, 0)$. Substituting (C.8) and (C.11) into the Euler equation and applying the same linearization arguments as in Proposition 4 yields

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

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. Expression (C.12) is the natural multi–factor generalization of (5.6): each physical, policy, or technology factor $f_{k,t+1}$, as well as their relevant quadratic combinations, is associated with a priced characteristic given by its exposure function in the dividend and damage processes.

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 exposure coefficients $(g_1(Z_{i,t}), \ldots, g_K(Z_{i,t}), \tilde{g}_{k\ell}(Z_{i,t}))$, and these coefficients 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. 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 5 and 4.6 are therefore robust to quadratic damages and to multiple climate shocks: the priced objects are exposure functions $g_k(Z_{i,t})$ determined by technology, policy, and preferences, and these need not collapse to a single ratio such as $X_{i,t}/V_{i,t}$.