What drives the greenium, and does it align with green innovations?

Nam Nguyen *

June 20, 2026

Abstract

The greenium—the expected-return spread between brown and green equity—is robustly estimated to be positive, yet two questions remain: what drives it, and does the higher cost of capital it imposes on brown firms spur the green innovation it is meant to encourage? This paper develops a continuous-time general-equilibrium model of a brown and a green sector that addresses both. The greenium decomposes into three interacting components: a *policy* component, from the brown sector's exposure to climate-policy tightening; a *diversification* component, because the dominant brown sector co-moves most with aggregate wealth and is the poorer hedge for the market portfolio; and a *hedging* component that runs the other way, as a household with a taste for green consumption fears green scarcity and accepts a lower return on green equity, masking part of the spread. Policy risk and diversification are the primary drivers, and the greenium deepens as climate risk intensifies. Brown firms can innovate by mitigating—lowering the frequency of policy shocks—or transitioning—converting brown capital into green—yet both are misaligned with the premium they face. This yields a novel, risk-based explanation for the recent decline in green innovation. Calibrated, the model produces a greenium of roughly 60 basis points at a brown share near three-quarters, rising to 75–85 as the brown sector dominates, in line with recent estimates.

Keywords: Green premium, Green consumption, ESG investing, Green innovation.

JEL codes: G11, Q50, Q51, Q55.

*The author is affiliated with Cambridge Endowment for Research in Finance, the University of Cambridge, email: nguyenn@jbs.cam.ac.uk.

1 Introduction

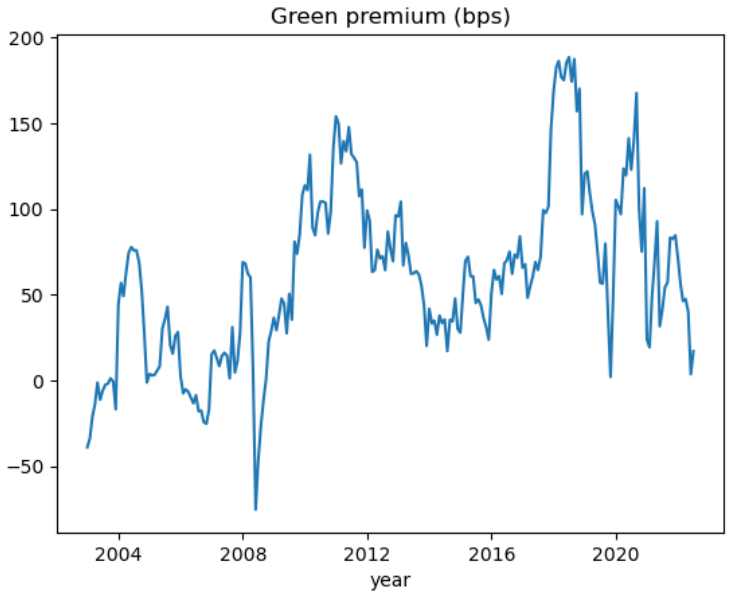
Climate change has put the pricing of environmental risk at the center of asset pricing. If investors demand a higher expected return to hold the equity of polluting firms, the resulting wedge in the cost of capital—the *greenium*, the brown-minus-green expected-return spread—is at once a gauge of how markets perceive climate risk and, in principle, a lever that should tilt capital and innovation toward the green economy. Measuring it has been contentious: realized returns send mixed signals, in good part because a decade of institutional reallocation toward green assets has repriced climate risk and pulled realized returns away from their expected counterparts. Forward-looking measures cut through this noise and point consistently to a positive premium—roughly 48–144 basis points in Pastor et al. (2022), 50–100 in Sautner et al. (2023b), and about 50 in Eskildsen et al. (2024), who further document that it has deepened over time. Figure 1 plots this brown-minus-green spread, which is both positive and widening.

A positive greenium poses two questions that the literature has largely addressed in isolation. The first is what *drives* it. The leading explanations—a non-pecuniary taste for holding green assets, the pricing of physical and transition risk, a portfolio-diversification motive, and a preference for green consumption—are typically modeled one at a time, leaving open how they interact and how large each is inside a single, internally consistent economy. The second is whether the greenium accomplishes what it is meant to. A higher cost of capital should prod brown firms to abate and to convert their operations to green; yet green innovation in the United States has been *falling* since the Paris Agreement, even as regulation has tightened and the greenium itself has risen (Figure 2). Whether and how the priced greenium feeds back into firms' real green-innovation decisions—and why that feedback might break down—has, to the best of my knowledge, not been examined in an equilibrium setting.

This paper develops a unified continuous-time general-equilibrium model that takes on both questions. A representative household with recursive preferences and a taste for green consumption holds claims on two sectors, green and brown, that differ in their exposure to climate-policy risk: brown output is intermittently stranded when policy tightens, at a frequency that rises with the brown sector's share of the economy. Crucially, brown firms are active rather than passive. Each can undertake two distinct green investments—*mitigation*, which lowers the frequency of policy shocks, and *transition*, which converts brown capital into green—so that the greenium and green innovation

Figure 1: **Green premium.**

Green premium. The following figure plots the expected return of the brown-minus-green factor from 2003 to 2023. The factor is a market-value-weighted long-short portfolio of U.S equities. By the end of each year t, firms are sorted within industries into terciles based on both their total and intensity of Scope 1 and Scope 2 emissions. The emission data is from Trucost and is lagged 12 months following Zhang (2024). Then, the terciles are combined across industries. Firms that fall in the top tercile based on both criteria are classified as brown, while those that belong to the bottom tercile are green. The difference in expected returns between brown and green firms is computed each month from July of year t+1 to June of year t+2. The expected return is the average of four implied cost of capital measures from Gebhardt et al. (2001), Claus and Thomas (2001), Easton (2004), and Ohlson and Juettner-Nauroth (2005). The data is provided by Eskildsen et al. (2024).

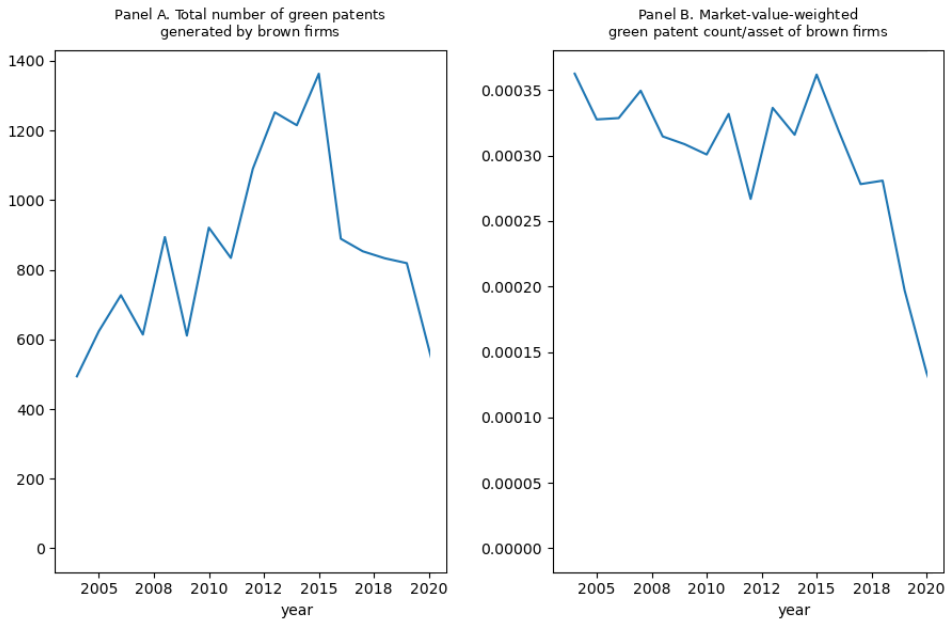


are determined jointly in equilibrium. The economy collapses to a single state variable, the brown output share η, which renders the model essentially closed-form and lets me trace every equilibrium object—the prices of risk, the greenium and its components, and the two investment policies—along the economy's path from brown to green. The analysis yields three main contributions.

First, the model delivers a transparent decomposition of the greenium into three economically distinct, mutually interacting components (Proposition 4 in Section 4.2). The *policy* component compensates investors for the brown sector's exposure to policy tightening; it is the largest of the three and climbs steeply as brown comes to dominate, both because shocks grow more frequent and because they arrive in states where the household's marginal utility is most sensitive. The *diversification* component arises because the larger sector co-moves more with aggregate wealth and is therefore the poorer hedge for the market portfolio: when brown dominates it is the worse

Figure 2: **Green innovation intensity.**

The following figure depicts the performance of U.S brown firms in generating green patents from 2004 to 2023. Panel A shows the total of green patents produced. Panel B plots the weighted average of patent count intensity as measured by the total number of patents/asset ratio. The updated data is provided by Kogan, Papanikolaou, Seru, and Stoffman (KPSS) following their work Kogan et al. (2017). The patents are categorized as "green" following the European Patent Office's Y02/Y04S tagging scheme.



diversifier and must offer a higher return, widening the spread. The *hedging* component runs in the opposite direction and partly masks the other two: because the household values green consumption and dreads green scarcity, green equity—which pays off precisely when green output is abundant—is a poor hedge against the costly scarce-green state, so investors accept a lower return to hold it, narrowing the brown-green gap. Policy risk and diversification are the primary drivers, while the hedging channel is second-order but not negligible. Calibrated to illustrative values, the model produces a greenium in the range of recent robust estimates—on the order of 60 basis points at a brown share near three-quarters and 75–85 basis points once the brown sector dominates—and it reproduces the deepening of the greenium over time as climate risk intensifies. The decomposition further shows that it is not only the average severity of a policy shock but also its *dispersion* that shapes the policy component.

Second—and this is the paper's central result—the model explains why a rising greenium has gone hand in hand with *falling* green innovation. Mitigation is the instrument that bears directly

on the priced policy risk, yet the two turn out to be sharply *misaligned*: mitigation is strongest while the brown sector is a minority and the policy premium is mild, and it shuts off exactly where that premium is largest. The mechanism is a general-equilibrium effect transmitted through the stochastic discount factor, and it hinges on a simple fact—mitigation can make policy shocks rarer but not milder. While brown is small, a shock destroys a little brown output but barely moves aggregate consumption, so it is a near-idiosyncratic loss the firm is happy to make less frequent. Once brown dominates, the same shock collapses *aggregate* consumption and sends the household's marginal utility soaring; because the household is risk-averse, that marginal-utility spike more than offsets the output destroyed, and a unit of brown payoff delivered in the disaster becomes worth at least as much to the household as one delivered in normal times. Reducing the frequency of such a disaster then adds nothing to firm value, and mitigation optimally falls to zero—in precisely the high-premium states where intuition would expect it to peak. The firms facing the steepest policy premium abate the least. This furnishes a novel, risk-based account of the decline in green innovation, distinct from the familiar appeal to uncertain or meager payoffs: tighter policy can be self-defeating at the margin, an unintended consequence of regulation working through equilibrium asset prices.

Third, the model characterizes the greenium and the two investments as a single coupled system (Section 5). The two priced risks map *asymmetrically* onto the two instruments. Mitigation enters the economy only through the policy-shock intensity, so it can compress the policy component but can never touch the diffusion risk behind the diversification component; the only way to shed that risk is to stop holding the brown asset, which is transition. Policy risk is thus reachable by either instrument, diversification risk by transition alone—and in the limit with no policy shocks, mitigation is zero everywhere while transition remains positive, sustained by the residual diversification gap. Transition itself tracks the valuation gap between green and brown equity—the price-level image of the risk greenium—until the green knowledge base it draws on thins out and caps it. The two investments, moreover, push the brown share in opposite directions: transition accelerates the greening of the economy, while mitigation, by thinning the very shocks that green it, brakes that process. Because the two substitute (mitigation props up the brown valuation and crowds transition out) while transition is self-correcting (greening erodes the premium that motivated it), the brown share comes to rest where these forces balance rather than running away.

This paper connects several strands of literature that have, for the most part, developed separately. It speaks, first, to the measurement of the greenium. Realized returns have produced conflicting verdicts—some studies find that brown equities outperform, others the reverse[1]—a disagreement that Ardia et al. (2023) and Pastor et al. (2022) trace to a repricing episode in which green outperformance reflects rising demand rather than lower expected returns. Forward-looking measures instead recover a robustly positive expected greenium (Pastor et al., 2022; Sautner et al., 2023b; Eskildsen et al., 2024). I take this positive, time-varying greenium as the object to be explained and supply the structural decomposition that the reduced-form estimates leave open.

Second, the paper contributes to the debate over the *source* of the greenium. One strand traces it to a non-pecuniary taste for holding green assets (Pedersen et al., 2021; Pastor et al., 2021, 2022; Zerbib, 2022). A recent survey of equity fund managers by Edmans et al. (2025), however, finds that financial considerations—downside protection, return enhancement, and volatility reduction—dominate stock selection on environmental grounds, echoing Krueger et al. (2020), who report that institutional investors regard climate risk, and regulatory risk in particular, as financially material and already materializing. A second, risk-based strand prices the underlying climate risks directly: Bansal et al. (2016) tie a temperature risk premium to long-run output damages, Sautner et al. (2023b) back a policy-uncertainty premium out of option prices, and Bolton and Kacperczyk (2023) show that transition risk pervades equity valuations, with firm-level exposure quantified by Sautner et al. (2023a). My framework belongs to this risk-based tradition and is consistent with the survey evidence, but instead of isolating a single channel it nests policy risk, diversification, and a consumption-based hedging motive together and quantifies how they interact.

Third, the paper builds on two channels that recent work has emphasized but not combined with policy risk. The diversification motive—that the dominant brown sector drives the volatility of the market portfolio, so green assets are prized for diversification—surfaces in the survey evidence of Edmans et al. (2025) and Krueger et al. (2020) and in the theory of Zerbib (2022) and Hambel et al. (2024), who show that risk mitigation and diversification are complements in the short run but trade off in the long run. The consumption channel—that a household taste for green goods makes green equity risky against future green scarcity—is identified empirically by Meier et al.

[1]Positive realized brown-minus-green spreads are reported by Bolton and Kacperczyk (2023, 2021); Hsu et al. (2023); the opposite findings appear in Pastor et al. (2022); Giese et al. (2021); Ardia et al. (2023); Bauer et al. (2022); Zhang (2024); Karolyi et al. (2023).

(2023) and modeled by Sauzet and Zerbib (2024), while Aghion et al. (2023) show that consumers' environmental preferences can steer firms toward clean innovation. I embed both channels in one equilibrium and show how each maps onto a component of the greenium and how they pull against one another.

Fourth, the paper relates to work on green innovation and the real effects of sustainable investing. A number of studies question whether investor pressure changes firm behavior in a material way (Berk and van Binsbergen, 2025; Heath et al., 2023; Kolbel et al., 2020), in contrast to the macroeconomic effect of carbon pricing documented by Colmer et al. (2024). On the value of abatement, Dietz et al. (2018) show that emissions-abatement projects can carry net present values that *rise* with their climate beta, and Hong et al. (2023a,b) show that mitigation postpones climate tipping points, raises welfare, and lowers firms' required returns, while Leippold and Yu (2025) find that innovative firms are more resilient to regulatory shocks; consistent with brown firms innovating more than green ones, Cohen et al. (2026) and Hartzmark and Shue (2023) document higher innovation among polluters. Against this backdrop, the post-2015 decline in U.S. green innovation (Figure 2) is a puzzle reminiscent of the green paradox. The paper's contribution is to rationalize it within the very equilibrium that prices the greenium: the misalignment between mitigation and the policy premium is a general-equilibrium consequence of how disasters are priced, not a breakdown of incentives at the level of the individual firm.

Finally, the paper builds methodologically on the production- and endowment-economy asset-pricing tradition with recursive preferences and multiple "trees" (Duffie and Skiadas, 1994; Cochrane et al., 2008). The policy-shock intensity that rises with the brown share follows Nguyen et al. (2024); the random shock size adapts the stochastic-recovery disaster of Hong et al. (2023a); and the transition technology is a knowledge-production function in the sense of Jones (1995, 2005), with a "standing-on-shoulders" spillover that embodies the directed-technical-change mechanism of Acemoglu et al. (2012)—who also stresses that the elasticity of substitution between clean and dirty goods is as consequential as the discount rate (here that elasticity exceeds one, in line with Papageorgiou et al. (2017); cf. Yin (2024)). Bringing these ingredients together is what lets a single tractable model produce both the greenium decomposition and the greenium-innovation feedback.

The remainder of the paper proceeds as follows. Section 2 sets out the two-sector economy, the policy-risk technology, the brown firm's mitigation and transition investments, and the household's

preferences. Section 3 states the household's and the brown firm's problems, characterizes the competitive equilibrium and the stochastic discount factor, and derives the optimal investment policies. Section 4 presents the asset-pricing implications and the greenium decomposition. Section 5 analyzes the interaction between the greenium and green innovation. Section 6 concludes. Proofs and derivations are collected in the Appendix.

2 Model

2.1 Endowments

The setting being considered is a continuous-time endowment economy with an infinite horizon. The economy is populated with a large number of firms classified into two types or sectors: eco-friendly green firms (the green sector) and non-eco-friendly brown firms (the brown sector). I denote by $Y_{n,t}$ the output of a type-n firm, where $n \in \{b, g\}$ with b standing for brown and g standing for green. The output processes are

$$\frac{dY_{b,t}}{Y_{b,t_-}} = \left(\mu_b - z_{m,t_-} - z_{t,t_-} - \psi_{t_-}\right) dt + \sigma_b dB_{b,t} + (Z-1)\, dN_{p,t}, \tag{2.1}$$

$$\frac{dY_{g,t}}{Y_{g,t_-}} = \left(\mu_g + \psi_{t_-} \frac{\eta_{t_-}}{1-\eta_{t_-}}\right) dt + \sigma_g dB_{g,t}. \tag{2.2}$$

$N_{p,t}$ represents the transitional policy risk, which follows a jump process with intensity λ_t. I assume the Brownian shocks $B_{b,t}$ and $B_{g,t}$ are uncorrelated with the jump process $N_{p,t}$, and $cov\left(dB_{b,t}, dB_{g,t}\right) = \nu\, dt$. All stochastic processes are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The brown firm controls a mitigation investment z_m and a transition investment z_t, and is subject to policy shocks of random size Z that arrive at the state-dependent rate λ_t; the transition investment induces an idiosyncratic transition rate ψ. I describe each object in turn; the brown firm's optimal choice of z_m and z_t is derived in Section 3.2.

Policy intensity. Following Nguyen et al. (2024), the number of policy-risk events follows a Poisson process whose instantaneous arrival rate increases in the fraction of aggregate output produced by the brown sector,

$$\lambda_t = \lambda_g + (\lambda_b - \lambda_g)\, \eta_t = \lambda_g + \eta_t \Delta_\lambda, \tag{2.3}$$

where $0 \leq \lambda_g \leq \lambda_b$ are two constants, $\Delta_\lambda = \lambda_b - \lambda_g$, and $\eta_t = \frac{Y_{b,t}}{Y_{b,t}+Y_{g,t}}$ is the share of brown output. The more predominant the brown sector, the larger the negative climate externality, which pushes Governments to adopt tighter environmental policies that disproportionately affect brown firms.

Random policy-shock size. When a policy shock materializes, brown output is scaled by a random factor $Z > 0$, so that $Y_{b,t} = Z\, Y_{b,t_-}$; this is the $(Z-1)\, dN_{p,t}$ term in (2.1). Modeling the shock size as a random variable, rather than as a fixed loss, follows Hong et al. (2023a), who let a Poisson disaster scale the capital stock by a stochastic recovery fraction whose distribution firms take as given. I let the regulatory response be two-sided. With probability p, the policy tightens, destroying a fraction $1 - Z^-$ of brown output, where the recovery fraction Z^- follows a Beta distribution supported on $(0, 1)$. With probability $1 - p$ the policy loosens, scaling brown output up by a factor Z^+ supported on $(1, 1 + \epsilon)$, where $\epsilon > 0$ caps the maximal relief. Z^+ follows a reversed-power distribution whose density is decreasing on the interval, so that mild relief (near $Z^+ = 1$) is more likely than large relief (near $1 + \epsilon$). The brown firm and the representative household take the resulting distribution Ξ of Z as given, with density

$$\xi(z) = \begin{cases} p\, \dfrac{z^{\beta-1}\,(1-z)^{a-1}}{B\,(\beta, a)}, & 0 < z < 1 \quad \text{(tightening)}, \\[2ex] (1-p)\, \dfrac{\alpha_r}{\epsilon^{\alpha_r}}\,(1+\epsilon-z)^{\alpha_r-1}, & 1 < z < 1+\epsilon \quad \text{(loosening)}, \end{cases} \tag{2.4}$$

where $\beta, a > 0$ govern the tightening severity, $\alpha_r \geq 1$ governs the loosening shape, and $B\,(\beta, a)$ is the Beta function. The firm and household form expectations $\mathbb{E}_Z[\cdot]$ over Ξ. The shock is summarized by its mean recovery $\mathbb{E}_Z[Z] = p\, \dfrac{\beta}{\beta + a} + (1-p)\left(1 + \dfrac{\epsilon}{\alpha_r + 1}\right)$ and its dispersion; the empirically relevant case is net expected tightening, $\mathbb{E}_Z[Z] \leq 1$. The purely destructive disaster of Hong et al. (2023a) is the special case $p = 1$ and $a = 1$ in which the loosening component carries no mass. As Section 4.2 shows, it is not only the mean of Z but also its dispersion that shapes the policy component of the green premium.

Mitigation. The brown firm can invest in mitigation (hereafter, green) technologies at intensity z_m to reduce its contribution to the climate externality and thereby lower the frequency of policy shocks that strand its assets:

$$\lambda_t \longrightarrow \frac{\lambda_t}{1 + \phi z_m}.$$

When a firm develops environmentally friendly technologies, it adapts better to the stricter standards that Governments require. The resource cost of mitigation is $z_{m,t^-}Y_{b,t^-}$, which is reflected in the drift of brown output in (2.1).

Transition. In addition to lowering the *frequency* of policy shocks through mitigation, the brown firm can invest in *transition* at intensity z_t to convert its operations into green ones. Transition arrives as an idiosyncratic Poisson event whose rate is increasing in the firm's transition investment and in the size of the green sector,

$$\psi\left(z_t,\eta\right)=\bar{\psi}\, z_t^{\beta_\psi}\left(1-\eta\right)^{\kappa}, \qquad \beta_\psi\in(0,1)\,,\ \kappa>0. \tag{2.5}$$

The specification (2.5) is a knowledge-production function in the sense of Jones (1995, 2005): the exponent $\beta_\psi<1$ captures diminishing returns to the firm's own transition effort (the "stepping-on-toes" effect), while $(1-\eta)^{\kappa}$ is a "standing-on-shoulders" knowledge spillover from the existing green stock. The larger the green sector ($1-\eta$ large), the more green technology, skilled labor, supply chains, and institutional know-how exist to build on, so transition is easier; conversely, when the green base is small ($\eta\to 1$) there is little to absorb and transition stalls. This is the directed-technical-change mechanism of Acemoglu et al. (2012), in which advances in a sector make subsequent advances in that sector more effective.

Upon a transition event, the firm's brown output is reclassified as green output. Because there is a continuum of brown firms, the law of large numbers implies that a flow $\psi\, Y_{b,t}\, dt$ of output is reclassified each instant. This produces the $-\psi$ term in the brown drift (2.1) and the corresponding inflow $+\psi\,\eta/(1-\eta)$ in the green drift (2.2), so that aggregate output $Y_{b,t}+Y_{g,t}$ is conserved by the reclassification. Brown shareholders receive the market value of the newly transitioned green equity as a spinoff, a dividend flow $\psi\, p_g\, Y_{b,t^-}\, dt$ that is made precise in Section 3.3. The resource cost of transition is $z_{t,t^-}Y_{b,t^-}$, also reflected in the drift of (2.1). Since transition pays off only when a green unit is worth more than a brown unit, $z_t=0$ (and hence $\psi=0$) whenever $p_g\le p_b$.

I assume that both the mitigation and transition investments are observable by investors.[2]

[2]For implications of non-observability on firms' green transitions, see Gupta and Starmans (2025).

2.2 Preferences

The representative consumer has a recursive preference over his/her consumption streams, i.e.

$$U_t = \mathbb{E}\left[\int_t^{\infty} f\left(C_s, U_s\right) ds\right], \tag{2.6}$$

where $f(\cdot)$ is the instantaneous Epstein–Duffie aggregator. I consider the special case of unit IES in which the aggregator simplifies to

$$f\left(C_t, U_t\right) = \rho(1-\gamma)U_t\left[log\left(C_t\right) - \frac{1}{1-\gamma} log\left((1-\gamma)U_t\right)\right],$$

in which $\rho > 0$ corresponds to the subjective discount factor, and γ is the degree of relative risk aversion. When $\gamma = 1$, the aggregator becomes $f\left(C_t, U_t\right) = \rho\left[log\left(C_t\right) - U_t\right]$. The representative consumer is endowed with claims on the output produced by both types of firms.

Following Sauzet and Zerbib (2024), the consumption bundle C_t is calculated using a modified Stiglitz–Dixit aggregator such that

$$C_t = \left[(1-\theta)^{\frac{1}{\alpha}} {C_{b,t}}^{\frac{\alpha-1}{\alpha}} + \theta^{\frac{1}{\alpha}} {C_{g,t}}^{\frac{\alpha-1}{\alpha}}\right]^{\frac{\alpha}{\alpha-1}}. \tag{2.7}$$

$C_{b,t}, C_{g,t}$ represent the consumptions of brown/green outputs at time t, and α controls the elasticity of substitution between them. This parameter α is, according to Acemoglu et al. (2012), as important as the discount rate when studying optimal responses to climate change. In addition, the reasonable assumption is $\alpha > 1$, consistent with Papageorgiou et al. (2017).[3]

The parameter $0 \leq \theta \leq 1$ controls altruistic preferences for the output of the green sector. Because the average investor prefers to consume green over brown output,[4] $\theta > \frac{1}{2}$. Note that in the limiting case where $\alpha \longrightarrow \infty$ (i.e., perfect substitution) and $\theta = \frac{1}{2}$ (i.e., no altruistic preference for green output), the consumption bundle becomes

$$C_t = C_{b,t} + C_{g,t}. \tag{2.8}$$

[3] Meanwhile, Yin (2024) shows that clean and dirty capital are complementary.

[4] Aghion et al. (2023) provide evidence that consumers' environmental preferences and product-market competition can jointly steer firms toward clean innovation, using panel data on automobile patents and measures of environmental willingness to pay.

3 Competitive Market Equilibrium

I focus on the decentralized equilibrium rather than the social planner's solution. The representative agent (the household) and the firms take equity and good prices as given and maximize lifetime utility and firm value, respectively. Goods and equity markets clear, $C_{b,t} = Y_{b,t}$, $C_{g,t} = Y_{g,t}$, and the representative agent holds all the dividend claims to both sectors. The Markovian Nash equilibrium of this stochastic differential game can be defined as in Carmona (2016). The state variables are $Y_{b,t}, Y_{g,t}, \lambda_t$, and W_t, where W_t is the total wealth of the representative agent. Following Cochrane et al. (2008) and Eberly and Wang (2025), I exploit the homogeneity of the model and reduce the dimension of the problem so that the only relevant state variable is η_t.

The evolution of the relative size of the brown sector, η_t, is given by

$$d\eta_t = \widehat{\Sigma}(\eta_{t_-})dt + \eta_{t_-}\left(1-\eta_{t_-}\right)\left[\sigma_b dB_{b,t} - \sigma_g dB_{g,t}\right] + \left[\eta_t^J(Z) - \eta_{t_-}\right] dN_{p,t}, \tag{3.1}$$

where the drift

$$\widehat{\Sigma}(\eta) = \eta\left(1-\eta\right)\left[\mu_b - z_m - z_t - \mu_g - \eta\sigma_b^2 + \left(1-\eta\right)\sigma_g^2 + \left(2\eta-1\right)\nu\sigma_b\sigma_g\right] - \psi\,\eta \tag{3.2}$$

contains the additional term $-\psi\,\eta$: because transition is an idiosyncratic Poisson event, the law of large numbers turns it into a deterministic outflow from the brown share, with no aggregate jump. The only jump in η comes from the aggregate policy shock. After each policy shock, the state moves discretely to

$$\eta_t^J(Z) = \frac{\eta_{t_-} Z}{\eta_{t_-} Z + \left(1-\eta_{t_-}\right)}, \tag{3.3}$$

which lies below η_{t_-} for a tightening shock ($Z < 1$) and above it for a loosening shock ($Z > 1$).

Let $m_{b,t}$ and $m_{g,t}$ be the prices of brown/green output. I normalize $\left[\left(1-\theta\right)m_{b,t}^{1-\alpha} + \theta\, m_{g,t}^{1-\alpha}\right]^{\frac{1}{1-\alpha}}$ to unity, so that the price of the aggregate consumption basket, $\frac{m_{b,t}C_{b,t}+m_{g,t}C_{g,t}}{C_t}$, equals one. Let $\xi_t = \frac{m_{g,t}}{m_{b,t}}$ be the relative price of green to brown output. The output prices are

$$m_{b,t} = \left[\left(1-\theta\right) + \theta\,\xi_t^{1-\alpha}\right]^{\frac{1}{\alpha-1}}, \qquad m_{g,t} = \left[\left(1-\theta\right)\xi_t^{\alpha-1} + \theta\right]^{\frac{1}{\alpha-1}}. \tag{3.4}$$

Using the homogeneity property, the equity price/output ratios of the two sectors are

$$p_{b,t} = \frac{P_{b,t}}{Y_{b,t}}, \qquad p_{g,t} = \frac{P_{g,t}}{Y_{g,t}}. \tag{3.5}$$

The equity returns of the brown and green firms are

$$dR_{b,t} = \Lambda_{b,t_-}\, dt + \Theta_{b,t_-}^T\, dB - \left[1 - \frac{Z\, p_b\left(\eta_t^J\right)}{p_b(\eta_{t_-})}\right] dN_{p,t}, \tag{3.6}$$

$$dR_{g,t} = \Lambda_{g,t_-}\, dt + \Theta_{g,t_-}^T\, dB - \left[1 - \frac{p_g\left(\eta_t^J\right)}{p_g(\eta_{t_-})}\right] dN_{p,t}, \tag{3.7}$$

where $dB = (dB_{b,t}, dB_{g,t})^T$. Two features distinguish the brown return. First, its jump carries the random shock size Z, which scales brown output at the policy event. Second, its drift Λ_b includes the spinoff term $\psi\,[p_g - p_b]^+ / p_b$: brown shareholders receive the market value of the newly transitioned green equity, net of the reclassified brown output. The green return is unaffected by Z (green output is not directly hit by the policy shock) and excludes the spinoff, which accrues to brown shareholders. The full expressions for $\Lambda_b, \Lambda_g, \Theta_b, \Theta_g$ are given in the Appendix.

The market equilibrium is characterized by the set of functions $\{p_b(\eta), p_g(\eta), \xi(\eta), r(\eta)\}$ representing the equity price/output ratios, the relative output price, and the instantaneous risk-free rate. To solve for the equilibrium, I employ the PDE method of Carmona (2016), which reduces to a system of Hamilton–Jacobi–Bellman equations derived from the brown firm's and household's problems, detailed in the Appendix.

3.1 Household's Problem

At any time t, the representative household takes the equilibrium variables as given and chooses consumption C_b, C_g and portfolio holdings π_b, π_g to maximize lifetime utility

$$V_t = \sup_{C_b, C_g, \pi_b, \pi_g} \mathbb{E}\left[\int_t^\infty f\left(C_s, V_s\right) ds\right], \tag{3.8}$$

subject to the budget constraint

$$\begin{aligned}\frac{dW_t}{W_{t_-}} &= \left[r_{t_-} - c_{b,t_-} m_{b,t_-} - c_{g,t_-} m_{g,t_-} + \pi_{b,t_-}\left(\Lambda_{b,t_-} - r_{t_-}\right) + \pi_{g,t_-}\left(\Lambda_{g,t_-} - r_{t_-}\right)\right] dt \\ &\quad + \left[\pi_{b,t_-}\Theta^T_{b,t_-} + \pi_{g,t_-}\Theta^T_{g,t_-}\right] dB \\ &\quad + \left\{\pi_{b,t_-} Z \frac{p_b\left(\eta^J_t\right)}{p_b(\eta_{t_-})} + \pi_{g,t_-}\frac{p_g\left(\eta^J_t\right)}{p_g(\eta_{t_-})} - \left(\pi_{b,t_-} + \pi_{g,t_-}\right)\right\} dN_{p,t}.\end{aligned} \tag{3.9}$$

As shown in the Appendix, the optimal consumption of brown/green output is

$$C_{b,t} = \frac{\rho W_t}{m_{b,t}\left[1 + \xi_t^{1-\alpha}\left(\frac{\theta}{1-\theta}\right)\right]}, \qquad C_{g,t} = \xi_t^{-\alpha}\left(\frac{\theta}{1-\theta}\right) C_{b,t}. \tag{3.10}$$

With unitary EIS, the price–dividend ratio of the aggregate consumption claim is constant, $\frac{m_{b,t}C_{b,t} + m_{g,t}C_{g,t}}{W_t} = \rho$, and, by the basket normalization,

$$C_t = \rho W_t. \tag{3.11}$$

Exploiting homogeneity, the household's value function takes the form

$$V_t = V\left(W_t, \eta_t\right) = \rho^{1-\gamma} G\left(\eta_t\right) \frac{W_t^{1-\gamma}}{1-\gamma}, \tag{3.12}$$

where the reduced-form G measures the (transformed) marginal value of wealth and satisfies the HJB equation given in the Appendix.

3.2 Brown Firm's Problem

At any time t, the brown firm takes prices and the dynamics of the stochastic discount factor (SDF) as given, and chooses its mitigation and transition investments z_m, z_t to maximize firm value,

$$P_{b,t} = p_b(\eta_t) Y_{b,t} = \sup_{z_m, z_t \geq 0} \mathbb{E}\left[\int_t^\infty \frac{M_s}{M_t}\left(m_{b,s} + \psi_s\, p_{g,s}\right) Y_{b,s}\, ds\right], \tag{3.13}$$

where the dividend flow $\left[m_b + \psi\, p_g\right] Y_b$ is the value of output sold at price m_b plus the spinoff value $\psi\, p_g Y_b$ of the green equity created by transition, and M_t is the equilibrium SDF.

Following Duffie and Skiadas (1994), $M_t = e^{\int_0^t f_V(C_s,V_s)ds} f_C(C_t, V_t)$, and no arbitrage implies

$$\frac{dM_t}{M_{t_-}} = -r_{t_-}dt - \Phi_{b,t_-}dB_{b,t} - \Phi_{g,t_-}dB_{g,t} + (\kappa_{p,t} - 1)\left(dN_{p,t} - \widetilde{\lambda}_{t_-}dt\right), \tag{3.14}$$

where $\widetilde{\lambda} = (\lambda_g + \eta\Delta_\lambda)/(1 + \phi z_m)$ is the mitigated policy intensity, and Φ_b, Φ_g and $\kappa_p - 1$ are the endogenous market prices of diffusion and policy risk, derived in the Appendix. κ_p measures the ratio of the SDF after versus before the shock, $M_t^J/M_{t_-} = \kappa_{p,t}$. Both the firm and the household price the policy jump under $\mathbb{E}_Z$.

As shown in the Appendix, the brown firm balances the resource cost of investing against two benefits—a lower exposure to policy-risk stranding through mitigation, and the spinoff gain from converting brown operations into more valuable green ones through transition:

$$\{z_m, z_t\} = \arg\max_{z_m, z_t \geq 0} \underbrace{p_b(\eta)(1 - z_m - z_t)}_{\text{Capital Gain}} + \underbrace{\left(\frac{\lambda_g + \eta\Delta_\lambda}{1 + \phi z_m}\right)\mathbb{E}_Z\left[\kappa_p Z\, p_b(\eta^J) - p_b(\eta)\right]}_{\text{Asset Stranding}} + \underbrace{\psi(z_t)\left[p_g(\eta) - p_b(\eta)\right]^+}_{\text{Transition Gain}}. \tag{3.15}$$

The two first-order conditions yield the optimal mitigation and transition investment intensities

$$z_m = \max\left\{0,\ \frac{1}{\phi}\left(\sqrt{\frac{\phi(\lambda_g + \eta\Delta_\lambda)}{p_b(\eta)}\left[\,p_b(\eta) - \mathbb{E}_Z(\kappa_p Z\, p_b(\eta^J))\,\right]^+} - 1\right)\right\}, \tag{3.16}$$

$$z_t = \left[\bar{\psi}\,\beta_\psi\,(1-\eta)^\kappa\,\frac{[p_g(\eta) - p_b(\eta)]^+}{p_b(\eta)}\right]^{\frac{1}{1-\beta_\psi}}. \tag{3.17}$$

Mitigation is governed by the stranding loss $-J \equiv p_b(\eta) - \mathbb{E}_Z\left[\kappa_p Z\, p_b(\eta^J)\right]$: the firm invests in mitigation only when this loss is positive, that is, when the fall in brown value at a policy shock is not offset by the accompanying rise in the SDF κ_p. Transition is governed by the value gap $[p_g - p_b]^+$, scaled by the knowledge spillover $(1-\eta)^\kappa$, and is active only when green equity is worth more than brown. The optimal investment behavior and the interaction between the two margins are analyzed in the sections that follow.

3.3 Characterization of the Market Equilibrium

I characterize the equilibrium variables $p_b(\eta), p_g(\eta), \xi(\eta)$ and $r(\eta)$, suppressing the dependence on t where possible. The closed-form relative output price is

$$\xi(\eta) = \left(\frac{\theta}{1-\theta}\right)^{\frac{1}{\alpha}}\left(\frac{\eta}{1-\eta}\right)^{\frac{1}{\alpha}}. \tag{3.18}$$

Figure 3: **Equilibrium relative price of the green/brown output.**

Equilibrium relative output price $\xi(\eta) = m_g/m_b$. Illustrative parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.

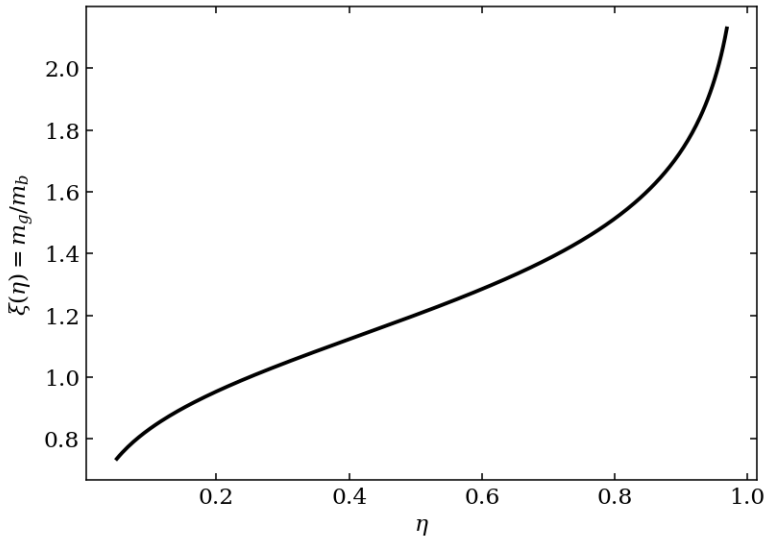


Figure 3 plots ξ. Because the household's CES basket makes it want to consume both goods, the marginal value of a good rises when its supply is limited. When the green sector is small (high η), green output is scarce and its price climbs relative to brown, raising ξ; in the low-η state, the reverse holds. Holding η fixed, ξ rises with the green-consumption bias θ and falls with the degree of substitution α: a stronger taste for green raises its price, while easier substitution dampens the household's sensitivity to scarcity.

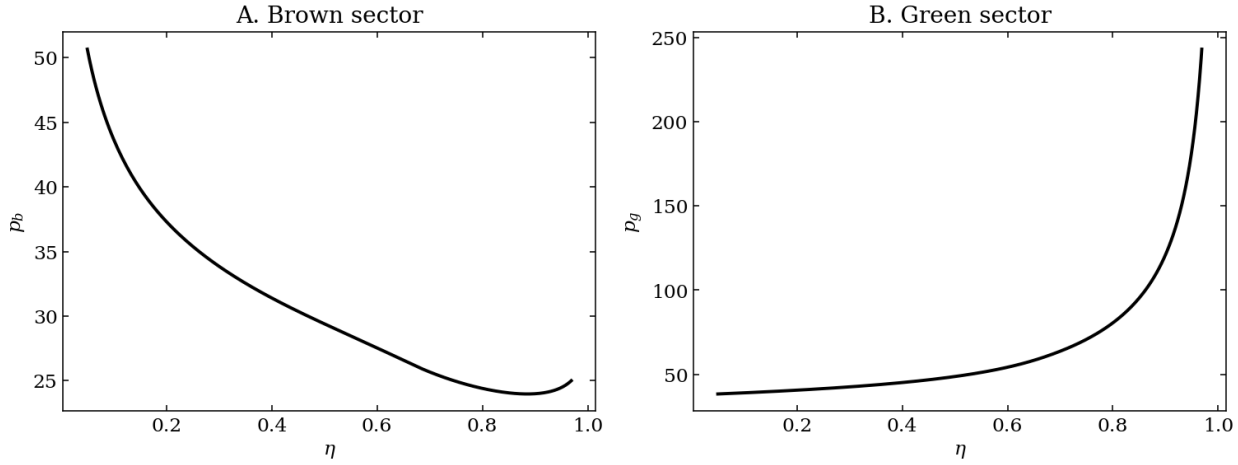
The equity valuation ratios show a stark contrast (Figure 4): p_b declines while p_g appreciates as the brown sector grows. Several forces shape them. First, the household wants a balanced portfolio to diversify sector risk; when one sector dominates, the other's equity co-moves less with wealth and is valuable for diversification, so the household pays more for it. Second, policy risk is positively priced and, in equilibrium, both equities load on it (Section 4.2), depressing valuations where shocks are frequent (high η). Third, the household's taste for green consumption raises p_g, especially when green output is scarce (high η), reinforcing the elevated right tail of p_g.

With the SDF dynamics (3.14) in hand, the equilibrium risk-free rate is characterized as follows.

Proposition 1. *The equilibrium interest rate has a diffusion-risk component reflecting precautionary saving against the Brownian shocks and a policy-risk component reflecting precautionary saving*

Figure 4: **Equilibrium valuation ratios.**

Equilibrium equity price/output ratios $p_b(\eta)$ (Panel A) and $p_g(\eta)$ (Panel B). Illustrative parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



against the jump,

$$r(\eta) = \underbrace{\rho + \mu_c(\eta) - \gamma\sigma_c^2(\eta)}_{\text{Diffusion Risks}} + \underbrace{\widetilde{\lambda}(\eta)\,\mathbb{E}_Z\left[\kappa_p(\eta, Z)\left(\pi_b(\eta)\frac{Z\,p_b(\eta^J)}{p_b(\eta)} + \pi_g(\eta)\frac{p_g(\eta^J)}{p_g(\eta)} - 1\right)\right]}_{\text{Policy Risk}}, \tag{3.19}$$

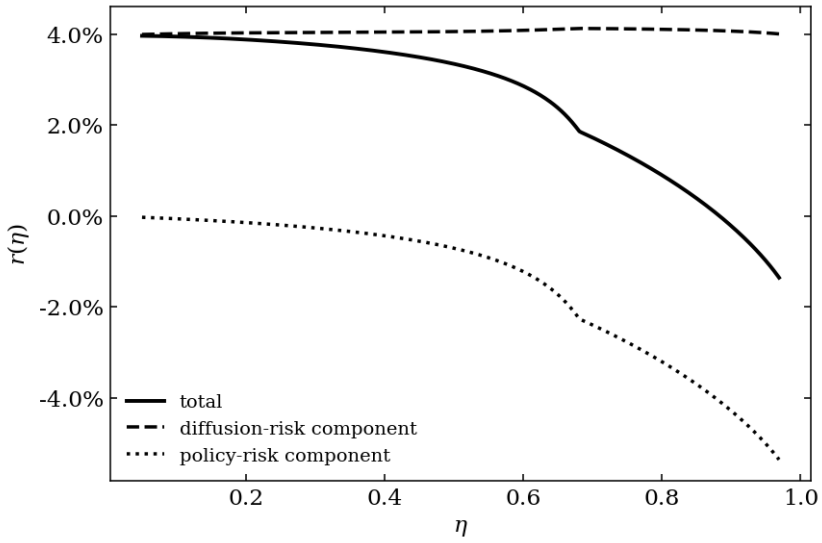
where μ_c and σ_c^2 are the drift and variance of optimal consumption growth, both given in the Appendix.

Equation (3.19) has the two components shown in Figure 5. The diffusion component (dashed) is hump-shaped: the diversification benefit is strongest when the two sectors' volatilities balance, which reduces total consumption volatility and the demand for the risk-free asset, pushing the rate down. The policy component (dotted) declines monotonically in η: the shock arrives more often and destroys more aggregate consumption as η rises, so precautionary demand for the safe asset strengthens and the rate falls, an effect attenuated by mitigation z_m.

The reduced-form value function G drives the marginal value of wealth and has a reverse-L shape with an elevated right tail (Figure 6). Absent the green-consumption taste and policy risk, G would be U-shaped because the CES basket makes the household dislike scarcity of either good, lifting the marginal value of wealth at both ends of η. The strong taste for green output raises

Figure 5: **Equilibrium risk-free rate.**

Equilibrium risk-free rate $r(\eta)$ (solid), with its diffusion-risk (dashed) and policy-risk (dotted) components from (3.19). Illustrative parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



that marginal value most where green is scarce (high η), since the household must then substitute toward the less-preferred brown good. In addition, policy risk, concentrated at high η, reinforces the right tail.

3.4 Solutions to the Household's and Brown Firm's Problems

3.4.1 The Household

Proposition 2. *The optimal risky portfolio of the representative household is*

$$\begin{bmatrix} \pi_b(\eta) \\ \pi_g(\eta) \end{bmatrix} = \frac{1}{\gamma}\Sigma(\eta)^{-1}\left\{ \begin{bmatrix} \Lambda_b - r(\eta) \\ \Lambda_g - r(\eta) \end{bmatrix} + \begin{bmatrix} \widetilde{\lambda}\,\mathbb{E}_Z\left[\kappa_p\left(\frac{Z\,p_b(\eta^J)}{p_b(\eta)} - 1\right)\right] \\ \widetilde{\lambda}\,\mathbb{E}_Z\left[\kappa_p\left(\frac{p_g(\eta^J)}{p_g(\eta)} - 1\right)\right] \end{bmatrix} + \Omega(\eta)\frac{G'(\eta)}{G(\eta)} \right\}, \tag{3.20}$$

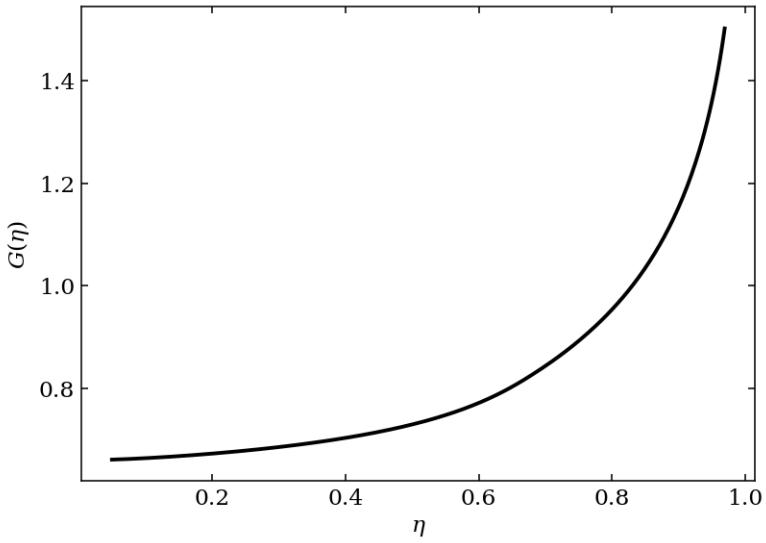
where $\widetilde{\lambda} = (\lambda_g + \eta\Delta_\lambda)/(1 + \phi z_m)$, Σ *is the return variance–covariance matrix, and* Ω *is the covariance vector between the returns and* η. *Both expressions are provided in the Appendix.*

Figure 7 shows the optimal risky portfolio across states. When the brown sector accounts for around 90 percent of total output, in line with recent estimates,[5] the household allocates

[5] See Clements et al. (2024).

Figure 6: **Household's reduced-form value function.**

Reduced-form value function $G(\eta)$. Illustrative parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.

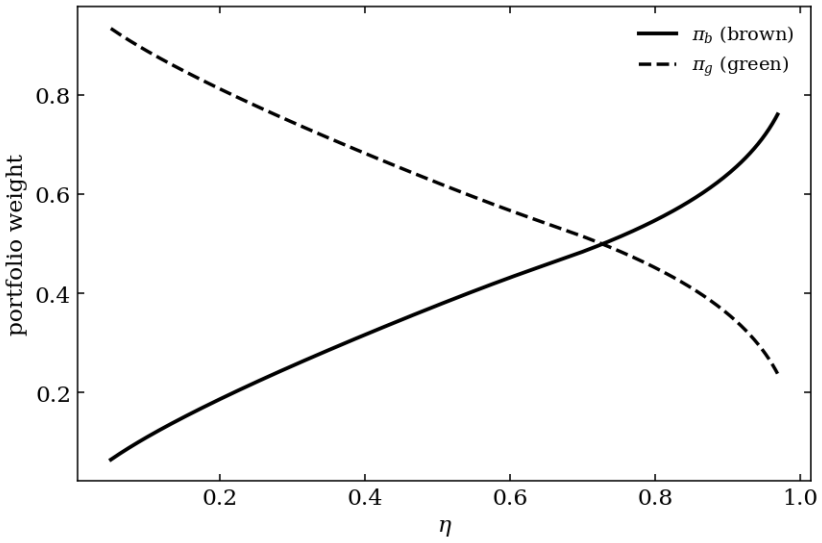


roughly 64 percent of its risky wealth to brown equity and 36 percent to green, broadly consistent with the evidence in Global Sustainable Investment Alliance (2020). As the economy transitions toward a greener composition, the household gradually prioritizes green equity and underweights polluting assets. A strand of the literature attributes such flows into green assets, and the associated expected-return spread, to a preference for green investing; here those patterns arise without any warm-glow investing channel.

Equation (3.20) satisfies a three-fund separation property. The first fund is the myopic mean–variance demand $\frac{1}{\gamma}\Sigma^{-1}(\Lambda - r)$, driven by the household's motive to diversify and reduce consumption volatility. The second is the tilt against policy risk: at a shock both equities lose value—brown output is scaled by Z, and green's valuation falls as η jumps left (Figure 4)—and the loss is amplified by the rise in the SDF κ_p, so the household tilts away from both risky assets. The magnitude of this tilt is attenuated by mitigation through $\widetilde{\lambda}$: if the policy intensity is abated, investors downweight policy risk in their portfolios. The third fund is the intertemporal hedge $\frac{1}{\gamma}\Sigma^{-1}\Omega\, G'/G$ against movements in η. Because the household has an altruistic taste for green consumption, it is averse to green scarcity (high η), where it must substitute toward the less-preferred brown good; it hedges by

Figure 7: **Optimal risky portfolio.**

Equilibrium portfolio weights $\pi_b(\eta)$ (brown) and $\pi_g(\eta)$ (green). Illustrative parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



underweighting green equity and overweighting brown, so that in aggregate the hedging components of the two assets offset and reallocate between them without changing total risky holdings. This hedge is why the equilibrium green holding sits below the level a myopic diversifier would choose.

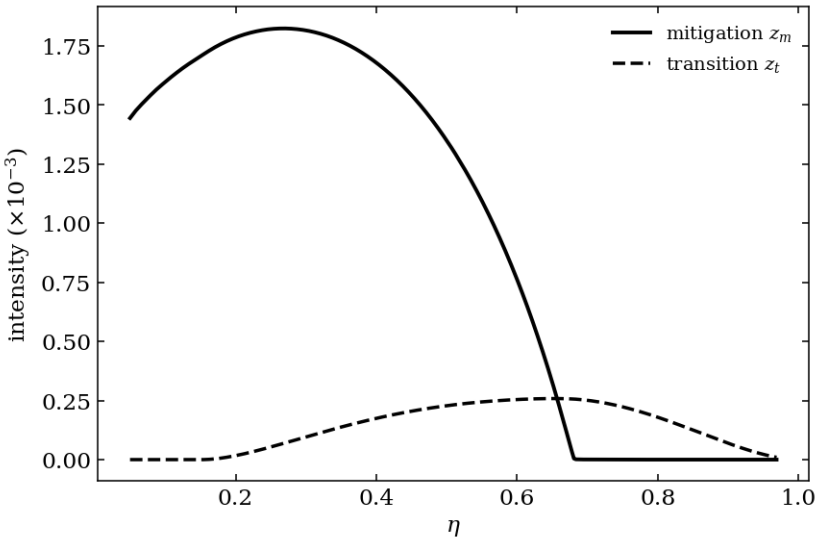
3.4.2 The Brown Firm

Figure 8 plots the brown firm's two investments. By the first-order conditions (3.16)–(3.17), the firm weighs the cost of forgone output against the benefit of each margin: mitigation lowers the frequency of policy shocks that strand its assets, while transition converts brown operations into more valuable green ones. When a policy shock materializes, the change in the brown firm's present value comes from three sources. First, the policy tightening scales down brown output by the factor Z; for instance, a carbon tax proportional to emissions shrinks the firm's cash flow. Second, the state η jumps to the left—by an amount that is largest when the two sectors are of comparable size—which revises the valuation ratio p_b. Third, because policy risk is positively priced, the SDF κ_p in (3.14) rises after the disaster-like event.

These three channels make *mitigation* hump-shaped, peaking at intermediate η and switching

Figure 8: **Brown firm's optimal investments.**

Optimal mitigation $z_m(\eta)$ (solid) and transition $z_t(\eta)$ (dashed) intensities, scaled by 10^3. Illustrative parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



off when the brown sector dominates. At the far right of η, the valuation ratio p_b falls only marginally at a shock, since η jumps just slightly to the left; and whatever drop occurs is more than masked by the rise in the SDF. The price of policy risk κ_p is large precisely when the brown sector dominates: in that region, shocks arrive frequently, and the repeated destruction of aggregate consumption renders each remaining unit of output highly valuable. This SDF effect raises the present value of each unit of brown output, and hence the firm's value to the investor, by more than the fall in the output level Y_b and in the ratio p_b. Counter to the intuition that a brown firm's value should collapse under intense policy tightening, the disaster effect actually *lifts* its present value. With no stranding loss left to insure against, the firm optimally sets $z_m = 0$ at high η.

At the opposite end, when η is very low, the price of policy risk is small, but the valuation ratio p_b rises sharply whenever the relative supply of brown output contracts—the extraordinary diversification benefit that scarce brown equity then provides to the household. This appreciation overrides much of the (now infrequent) loss of output at a shock, so the firm again gains little from curbing its policy exposure and invests only modestly in mitigation. Mitigation is therefore strongest at intermediate η, where the stranding loss is largest: the price of policy risk is high

enough to matter, while the offsetting SDF and diversification effects are not yet dominant.

Transition obeys a different logic. By (3.17), z_t is positive only where green equity is worth more than brown ($p_g > p_b$), and is increasing both in the value gap $[p_g - p_b]^+$ and in the knowledge spillover $(1-\eta)^\kappa$. As η rises, green output grows scarce, and its valuation climbs, widening the gap (Figure 4) and pulling transition up; but the spillover $(1-\eta)^\kappa$ shrinks as the green base vanishes, pulling it down. The net effect is again a hump: transition is muted at low η, where the gap is small or negative; it peaks at intermediate η; and it is throttled as $\eta \to 1$, where the thinning green sector leaves little technology, skilled labor, or supply-chain capacity to build on.

Comparing the two margins reveals a misalignment with the very risk they respond to. Mitigation is active where policy risk is moderate and switches off exactly where policy risk is most intense (high η), because the disaster effect there eliminates the firm's private incentive to abate. Transition takes over at intermediate η but is itself starved at high η by the thin green base. Thus, in the brown-dominated states where policy risk—and the externality it reflects—is greatest, neither private margin is strongly active: mitigation is futile, and transition is throttled. This tension between the size of the greenium and the green innovation it is meant to elicit is the focus of the next section.

4 Asset Pricing Implications

4.1 Prices and Quantities of Risk

Proposition 3. *From the stochastic discount factor (3.14), the equilibrium prices of diffusion risk Φ_b, Φ_g and of policy risk κ_p are*

$$\Phi_b(\eta) = \gamma\big[\Gamma_b\pi_b(\eta) + \Gamma_g\pi_g(\eta)\big] - \frac{G'(\eta)}{G(\eta)}\eta(1-\eta)\sigma_b, \tag{4.1}$$

$$\Phi_g(\eta) = \gamma\big[-\Delta_b\pi_b(\eta) + \Delta_g\pi_g(\eta)\big] + \frac{G'(\eta)}{G(\eta)}\eta(1-\eta)\sigma_g, \tag{4.2}$$

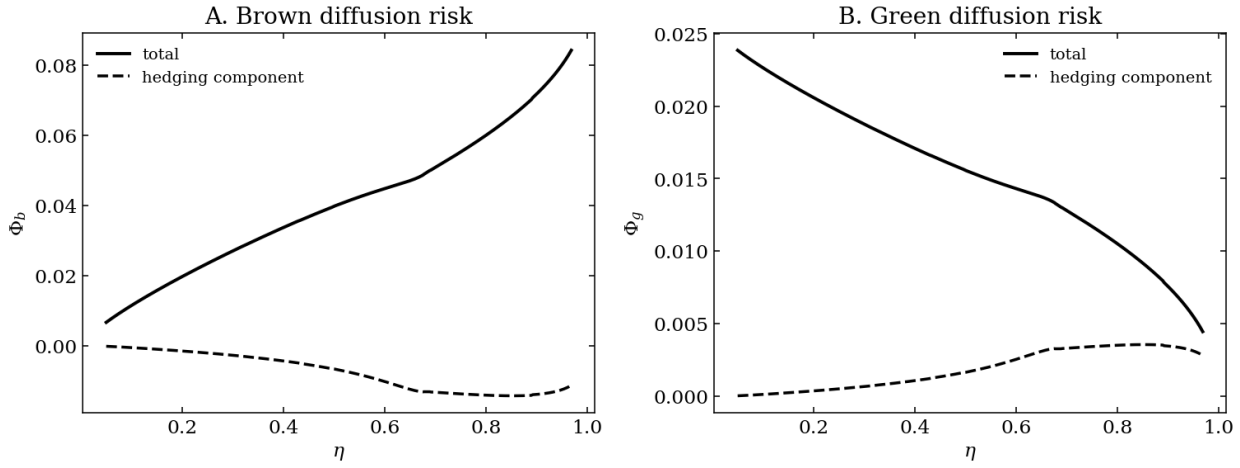
and

$$\kappa_p(\eta, Z) = \left(\pi_b(\eta)\frac{Z\,p_b(\eta^J)}{p_b(\eta)} + \pi_g(\eta)\frac{p_g(\eta^J)}{p_g(\eta)}\right)^{-\gamma}\frac{G(\eta^J)}{G(\eta)}. \tag{4.3}$$

The prices of diffusion risk in (4.1)–(4.2) each combine two elements. The first is a myopic market component, capturing the reward the household demands for holding assets that co-move

Figure 9: **Equilibrium prices of diffusion risk.**

Solid lines plot the total prices of diffusion risk $\Phi_b(\eta)$ and $\Phi_g(\eta)$; dashed lines plot their hedging components, so the gap between each solid line and its dashed counterpart is the market component. Parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



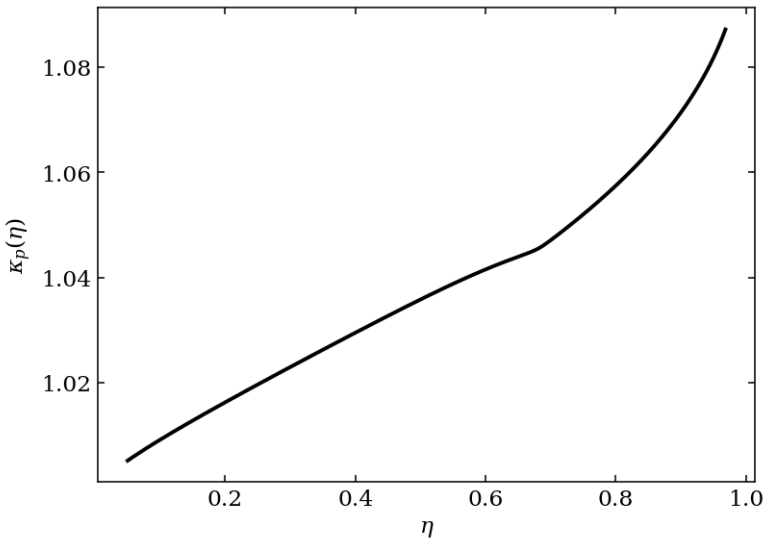
with its wealth. This component tracks the relative size of the sector: in Figure 9, it corresponds to the gap between each solid line and its dashed counterpart. When a sector is large, it makes up most of the market portfolio, a shock to its output drives a large share of total wealth volatility, and the household commands a high reward per unit of that risk; the smaller sector carries a low market price of risk. Accordingly, the market component of the brown price of risk rises with η while the green component falls.

The second element is the hedging component, which stems from the sensitivity of the marginal value of wealth to the state. Because the household's basket tilts toward green output ($\theta > \frac{1}{2}$), it is especially averse to states in which green is scarce. The green asset pays off precisely when green output is abundant and is therefore a poor hedge against that costly state, so the green-sector shock carries a positive hedging component. The brown good substitutes well for green when green is scarce, making the brown asset a good hedge, so its hedging component is negative. Figure 9 shows exactly this: the dashed green line sits above its solid line (positive hedging premium) while the dashed brown line sits below (negative hedging discount).

Figure 10 shows that the price of policy risk κ_p rises monotonically with the brown share. This term drives the risk-free rate, the equity risk premia, and the brown firm's optimal innovation. Since the effective price of the jump is $\kappa_p - 1$, comparison with the prices of diffusion risk shows it

Figure 10: **Equilibrium price of policy risk.**

The figure plots the modified price of policy risk $\kappa_p(\eta)$ averaged over the shock distribution. Parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



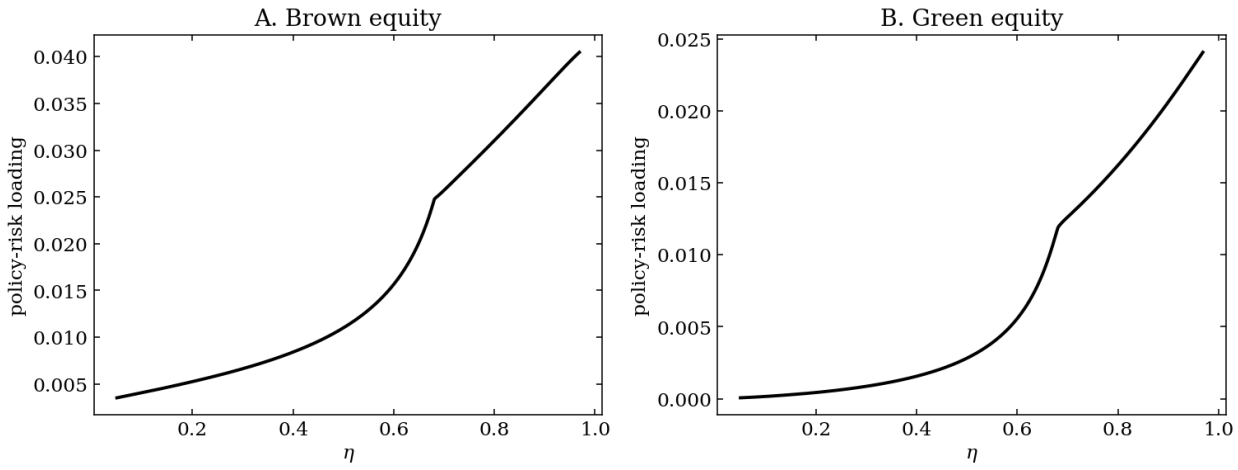
is sizeable. From (4.3) and (3.12), κ_p equals the ratio of the household's marginal value of wealth just after the shock to its value just before, so the household demands more compensation for the policy risk precisely when a shock raises its marginal value of wealth more sharply. The mild kink near $\eta \approx 0.74$ coincides with the point at which the brown firm ceases to invest in mitigation, discussed below.

The quantities of risk in Figures 11–12 display informative patterns. Figure 11 reports each asset's exposure to the policy shock. Both exposures are negative: when the shock lands, brown output is struck directly and the brown return falls, while the accompanying drop in η lowers the green valuation ratio p_g (which rises in η, as in Figure 4) and so drags the green return down as well. The green asset is thus exposed to the policy risk even though only the brown firm is directly targeted, a transmission that works entirely through the relative-size channel analyzed in Section 3.3. Brown remains the more exposed of the two at every η, and both exposures steepen as η rises, with the same mitigation-shutoff kink near $\eta \approx 0.74$.

Their exposure to the diffusion risks differs more sharply. Figure 12 plots the covariance of each return with the two Brownian shocks. Although the shocks are sector-specific, each spreads to the other sector's equity through its effect on the state variable, so both returns load on dB_b and

Figure 11: **Loadings of assets on the policy risk.**

The figure plots the policy-shock exposure $\mathbb{E}_Z[\cdot]$ of brown and green equity. Parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



dB_g.[6] The brown return is governed by its own shock throughout, the standard case. The green return is not: it covaries mainly with its own shock when the green sector is large, but increasingly with the brown shock once the brown sector dominates, the two curves crossing near $\eta \approx 0.47$. At that point, the green return is effectively driven by the state η, which is itself moved chiefly by the brown shock, so green equity inherits brown-sector risk exactly when the market price of that risk is high.

4.2 The Greenium

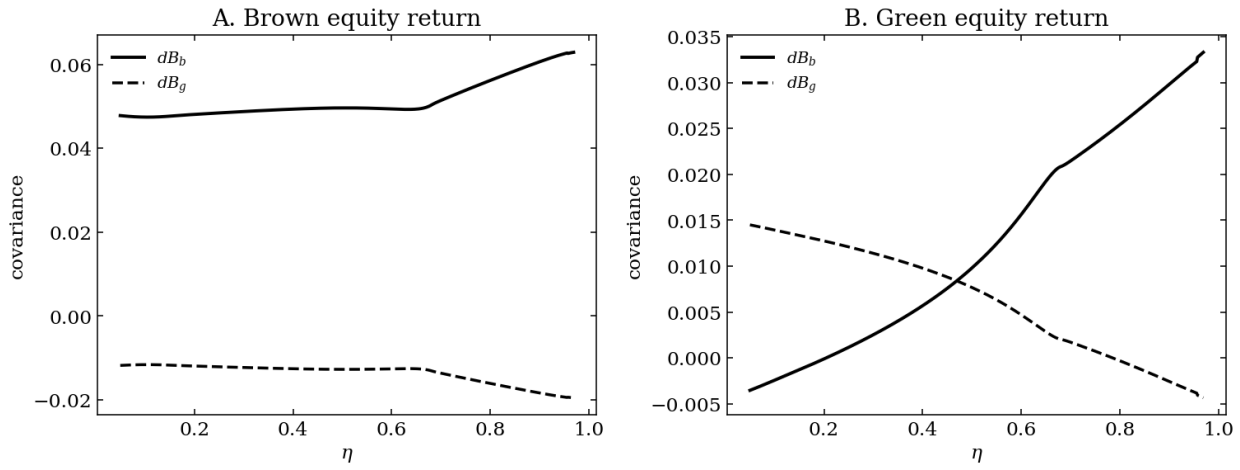
Proposition 4. *The greenium, defined as the spread between the expected excess returns on brown and green equity, decomposes into a diversification, a hedging, and a policy component,*

$$\underbrace{\big(\mathbb{E}[dR_b]/dt - r\big) - \big(\mathbb{E}[dR_g]/dt - r\big)}_{\text{Greenium}(\eta)} = \underbrace{\gamma \left\langle \sigma_{R_b} - \sigma_{R_g},\, \sigma_W \right\rangle_\nu}_{\text{Div}(\eta)} \underbrace{- \frac{G'(\eta)}{G(\eta)} \left\langle \sigma_{R_b} - \sigma_{R_g},\, \sigma_\eta \right\rangle_\nu}_{-\,\text{Hdg}(\eta)}$$
$$\underbrace{- \widetilde{\lambda}(\eta)\, \mathbb{E}_Z\Big[\big(\kappa_p - 1\big) \Big(\frac{Z\, p_b(\eta^J)}{p_b(\eta)} - \frac{p_g(\eta^J)}{p_g(\eta)} \Big) \Big]}_{\text{Pol}(\eta)}, \tag{4.4}$$

[6] The same cross-loading through a common state variable is documented by Cochrane et al. (2008).

Figure 12: **Loadings of assets on the diffusion risks.**

Each panel plots the covariance of an equity return with the two Brownian shocks dB_b and dB_g. Parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



where $\sigma_{R_b} = (\Gamma_b, -\Delta_b)$ *and* $\sigma_{R_g} = (\Gamma_g, \Delta_g)$ *are the return loadings,* $\sigma_W = \pi_b \sigma_{R_b} + \pi_g \sigma_{R_g}$, $\sigma_\eta = (\eta(1-\eta)\sigma_b, -\eta(1-\eta)\sigma_g)$, *and* $\langle u, v \rangle_\nu = u_1 v_1 + u_2 v_2 + \nu(u_1 v_2 + u_2 v_1)$.

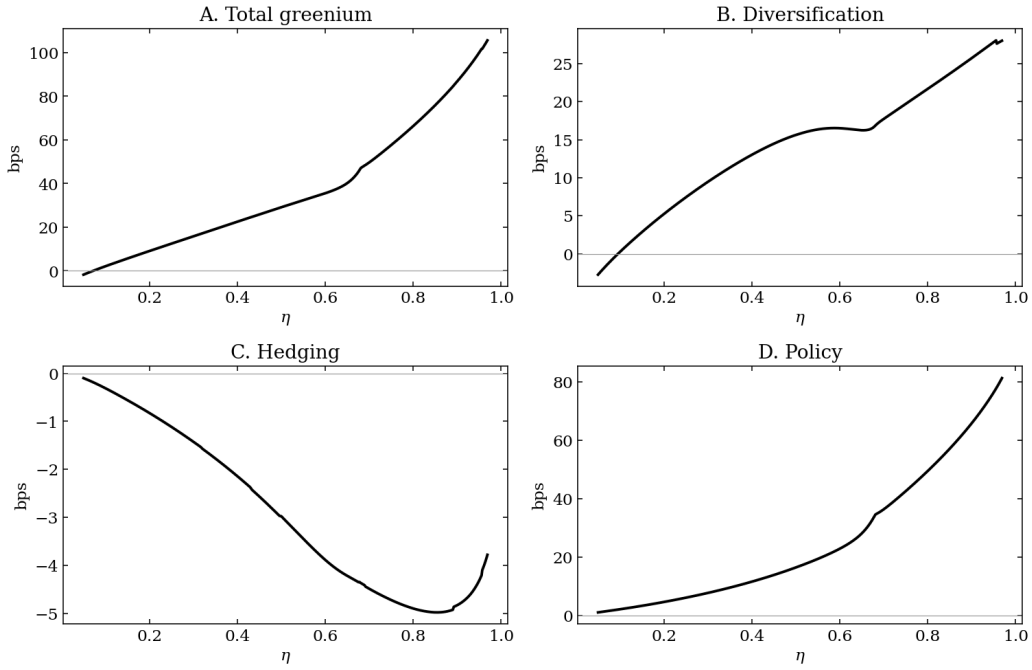
Figure 13 plots the greenium in Panel A and its diversification, hedging, and policy components in Panels B, C, and D. Panel A shows that the total spread rises steadily as the brown sector expands: at a brown output share of 85–90 percent it reaches roughly 75–85 basis points, broadly consistent with recent robust estimates of the greenium. The decomposition in (4.4) shows what drives it.

The diversification component reflects the relative hedging value of the two assets for the market portfolio. Whichever sector is larger co-moves more with total wealth and is the worse diversifier, so it must offer a higher premium on this account. As Panel B shows, the component is small and slightly negative when the green sector dominates at the far left, and grows along the increasing path of η as brown becomes the dominant, less-diversifying asset, reaching about 25 basis points at the high end.

The hedging component, plotted in Panel C, works in the opposite direction and slightly compresses the greenium. As the brown sector grows and green output becomes scarce, the household's marginal value of wealth rises in exactly the state where green equity pays off least, so green carries a hedging premium that narrows the brown−green spread. This component is

Figure 13: **Greenium decomposition.**

Panel A plots the total greenium (brown−green expected excess return); Panels B, C, and D plot its diversification, hedging, and policy components, all in basis points. Parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



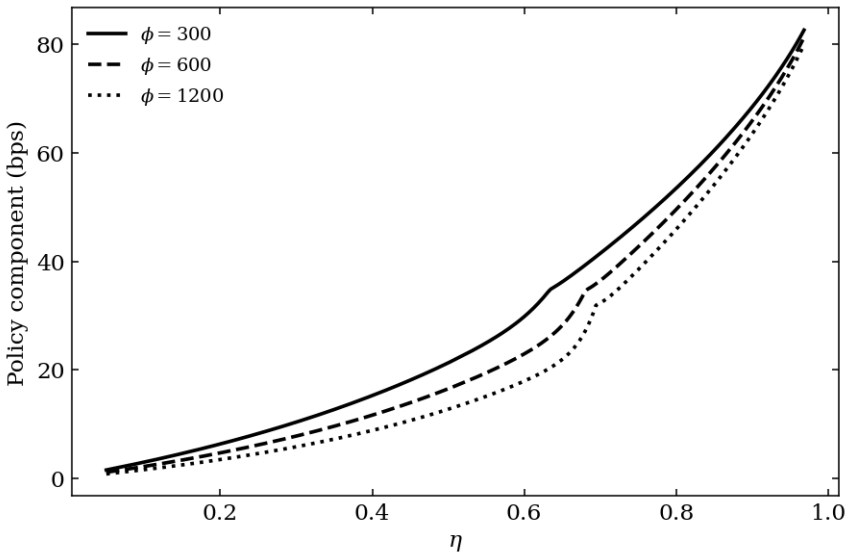
modestly negative throughout, of the order of 5 basis points when the green share falls to 10–15 percent.

The policy component, plotted in Panel D, is the largest contributor and accounts for the bulk of the greenium at high η. Because brown output is far more likely to be stranded when climate policy tightens, the household demands a substantial premium to hold brown equity. Equation (4.4) shows that this component widens both with the intensity of policy risk and with the brown firm's differential exposure to it, and that it scales with the price of policy risk κ_p; since both the intensity and κ_p increase in η, the policy component rises steeply as the brown sector expands. Its visible kink near $\eta \approx 0.74$ marks the brown share above which the firm stops investing in mitigation, beyond which the policy component is no longer attenuated and climbs most rapidly.

Mitigation by the brown firm scales down the intensity of policy risk and so counteracts the policy component. Its optimal intensity in Figure 8 is hump-shaped because the firm balances curbing the policy shock against the loss of current output, and it falls to zero once the brown

Figure 14: **Policy component under different abatement plans.**

The figure plots the policy component of the greenium for three values of the mitigation-effectiveness parameter ϕ; a higher ϕ means each unit of mitigation investment cuts the policy intensity by more. Other parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



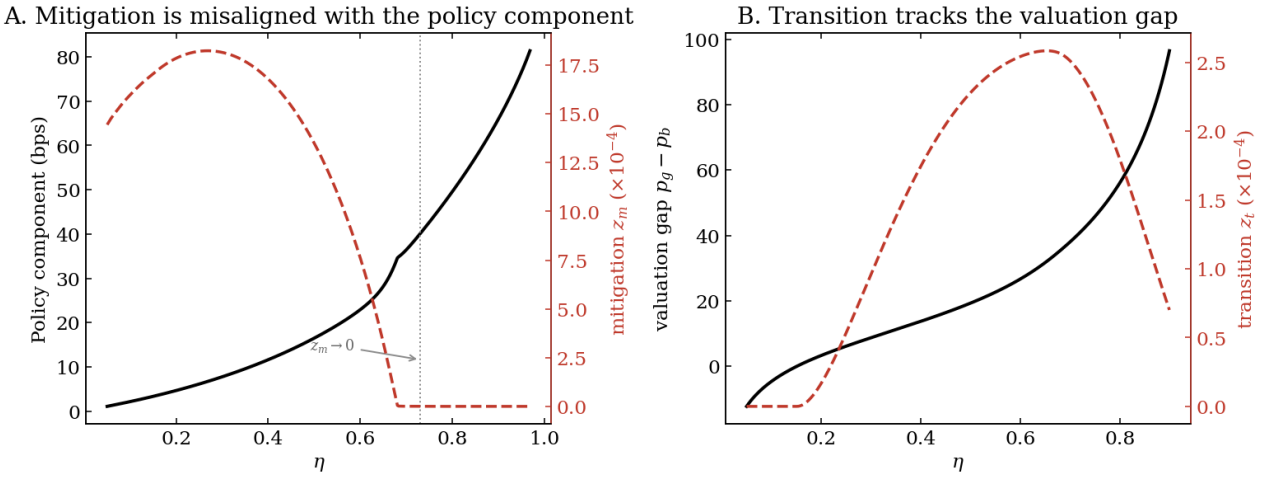
share is high enough that the disaster has become nearly unavoidable, as discussed in Section 3.4.2. Figure 14 traces the effect of more effective mitigation on the policy component. A higher ϕ lowers the component noticeably when η is small or moderate, but the three curves converge as η rises and become almost indistinguishable where the greenium is largest. The reason is structural: the firm abates when the arrival intensity is still mild, not when it is severe, so the shielding it provides is concentrated precisely where the policy component is already small. Together with the hedging motive against green scarcity, the firm's limited abatement leaves the greenium substantial, which helps reconcile the model with the empirical counterpart.

5 The Interaction between the Greenium and Green Innovation

The central question of the paper is how the components of the greenium interact with the firm's two channels of green innovation, mitigation, and transition. The brown firm faces the same policy and market risks that the household prices, and its two investments answer different facets of that risk. Recall the first-order conditions of Section 3.4.2. Mitigation responds to the

Figure 15: **The greenium and green investment.**

Panel A overlays the policy component of the greenium (left axis) and the optimal mitigation intensity z_m (right axis); the dotted line marks the brown share at which mitigation switches off. Panel B overlays the optimal transition intensity z_t (left axis) and the valuation gap $p_g - p_b$ (right axis). Parameter values are $\gamma = 1.7$, $\rho = 0.025$, $\mu_b = 0.02$, $\sigma_b = 0.06$, $\mu_g = 0.015$, $\sigma_g = 0.015$, $\nu = -0.3$, $\theta = 0.75$, $\alpha = 6$, $\lambda_g = 0.14$, $\Delta_\lambda = 0.75$, $\phi = 600$, $\bar{\psi} = 0.13$, $\beta_\psi = 0.5$, $\kappa = 1.5$, $\mathbb{E}[Z] = 0.955$.



stochastic-discount-factor-weighted capital loss a policy shock inflicts on brown equity,

$$-J \;=\; p_b - \mathbb{E}_Z\left[\kappa_p \, Z \, p_b^J\right], \tag{5.1}$$

and is positive only when $-J > 0$, that is, only when the loss is partly mitigable. Transition responds to the valuation gap $(p_g - p_b)^+$, gated by the knowledge spillover $(1-\eta)^\kappa$. The two are active in different regions of the state space (Figure 8): mitigation peaks while the brown sector is a minority, near $\eta \approx 0.27$, and transition peaks while it is larger, near $\eta \approx 0.65$. The firm therefore *defends* itself by abating while brown is small and *exits* into green as it grows — a handoff from mitigation to transition.

The policy component is the part of the greenium that mitigation acts on, yet the two are sharply *misaligned*. Panel A of Figure 15 overlays the policy component and the mitigation schedule. The policy component rises monotonically in η and is largest where the brown sector dominates; mitigation does the opposite, peaking in the interior and switching off near $\eta \approx 0.74$, exactly where the policy component is largest.

The reason is what mitigation can and cannot do. It lowers the *intensity* of the policy shock,

$\widetilde{\lambda}$ — how often the shock arrives — not the output Z it destroys when one occurs. Its marginal value is the shadow loss $-J = p_b - \mathbb{E}_Z[\kappa_p Z p_b^J]$: the value of brown equity if the shock is averted, p_b, net of its marginal-utility-weighted value if the shock lands, $\mathbb{E}_Z[\kappa_p Z p_b^J]$. Making the shock rarer is worthwhile only when a shock is, on net, a loss the household dislikes bearing — only when $-J > 0$. Two offsetting forces sit inside the second term: the shock destroys brown output ($Z < 1$), which lowers the payoff, while κ_p, the jump in the household's marginal utility, raises the value of whatever payoff is delivered in the shock state.

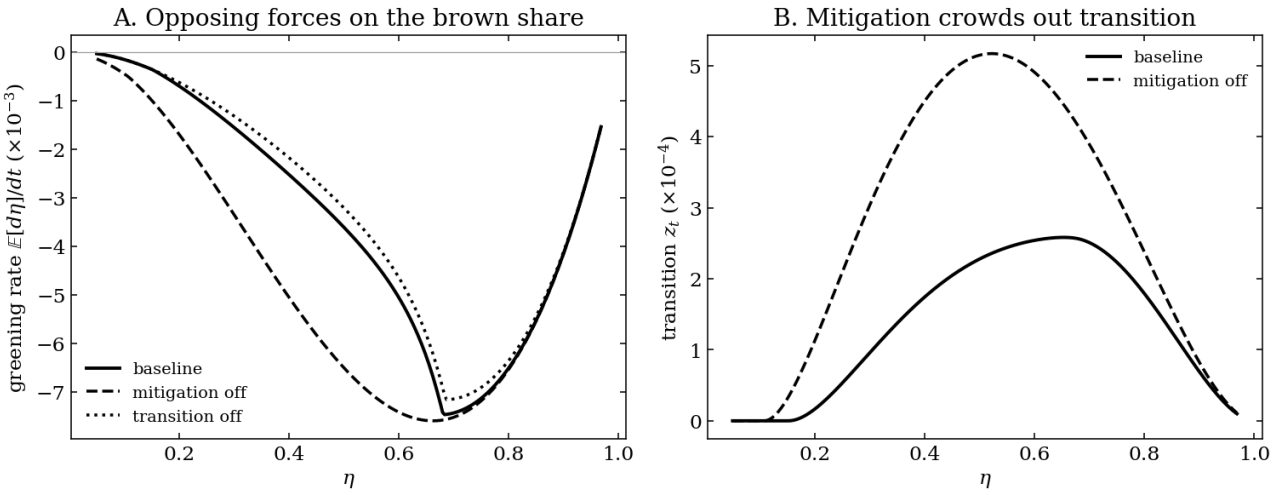
Which force dominates depends on the size of the brown sector, because that is what makes the shock idiosyncratic or aggregate. While brown is a minority, it is a small part of total output: a policy shock destroys brown capital but barely moves *aggregate* consumption, so the household's marginal utility is almost unchanged when the shock strikes ($\kappa_p \approx 1$), and the risk is partly mitigable. The loss is then close to a pure output loss felt at ordinary marginal utility, $\mathbb{E}_Z[\kappa_p Z p_b^J] < p_b$, so $-J > 0$ and the firm abates. Once brown is large, the shock turns *aggregate*: brown is essentially the whole economy, so the same shock collapses aggregate consumption — and it is that collapse that drives marginal utility up. The high marginal utility is therefore a *consequence* of the shock, not a coincidence of timing, and the spike is precisely κ_p, large now because the household holds the market and cannot diversify the loss away. Since the household is risk-averse ($\gamma > 1$), the proportional jump in marginal utility outweighs the proportional output loss, so in the household's valuation a unit of brown payoff delivered in the disaster is worth at least as much as one delivered in normal times: $\mathbb{E}_Z[\kappa_p Z p_b^J] \geq p_b$ and $-J \leq 0$.

The implication for mitigation is immediate. Lowering the intensity only shifts probability from the shock state to the no-shock state; when the shock state is, in marginal-utility terms, no less valuable than the no-shock state ($-J \leq 0$), making it rarer adds nothing to the firm's value, and mitigation falls to its lower bound of zero. Mitigation pays only while brown is a minority, and its policy risk is partly mitigable. Where it operates, it compresses the policy component materially (Figure 14); removing it would raise that component by roughly seventy percent at $\eta = 0.5$, but only in the interior, since at high η mitigation is already absent. The firms facing the largest policy premium abate the least.

Transition, in contrast, tracks the valuation gap. Panel B of Figure 15 shows the transition schedule rising with $p_g - p_b$, which is the price-level image of the risk greenium: the diversification

Figure 16: **Opposing dynamics and crowding out.**

Panel A plots the expected greening rate $\mathbb{E}[d\eta]/dt$ in the baseline economy and in two counterfactuals, one with mitigation switched off, and one with transition switched off. Panel B plots the optimal transition intensity z_t in the baseline economy and with mitigation switched off. Parameter values are as in Figure 15.



and policy components depress p_b and so widen the gap. Transition follows the gap upward through the interior; then the spillover $(1-\eta)^\kappa$ — the green knowledge base it can build on — shrinks and caps it even as the gap keeps widening, producing the hump in z_t.

The two risk-based components map *asymmetrically* onto the two investments, and this asymmetry is the heart of the interaction. Mitigation enters the model in exactly one place, the policy intensity $\widetilde{\lambda} = (\lambda_g + \eta\Delta_\lambda)/(1+\phi z_m)$; it scales the policy jump and nothing else, leaving the diffusion loadings (Γ_b, Δ_b) that generate the diversification component untouched. Mitigation can therefore compress the policy component but can never reduce the diversification component — the only way to shed diffusion risk is to stop holding the brown asset, which is transition. The policy component is thus addressable by *either* instrument, while the diversification component is addressable by transition *alone*. The cleanest confirmation is the limit with no policy shocks: setting $\lambda = 0$ leaves mitigation with nothing to act on, so $z_m = 0$ everywhere, while transition remains positive, sustained by the residual diversification gap. Mitigation is, by construction, a tool against the jump; transition is a wholesale exit that addresses the firm's entire risk profile. In magnitudes, the policy component dominates the gap, so the diversification component is the second driver of transition rather than the first; the point is the asymmetry in what each instrument *can* reach, not that diversification outweighs policy.

The two investments also pull the brown share in *opposite* directions. Transition converts brown capital to green and enters the drift of η directly through $-\psi(z_t)\eta$, accelerating the greening of the economy. Mitigation works the other way: by lowering the policy intensity, it thins the policy shocks that pull η down, so it *brakes* greening. Panel A of Figure 16 makes this concrete with two counterfactuals. Removing mitigation makes the brown share fall faster — the expected greening rate at $\eta = 0.5$ moves from -3.6 to -6.5×10^{-3} — while removing transition makes it fall slower, to -3.2×10^{-3}. Mitigation is a brown-preserving brake; transition is the greening accelerant.

The two instruments also substitute for one another. Mitigation props up p_b, which narrows the valuation gap and crowds out transition: with mitigation switched off, the gap widens and the transition schedule more than doubles (Panel B of Figure 16). And transition is self-correcting. By driving η down, it moves the economy into states where the policy component and the gap are both smaller, eroding the very premium that motivated it. The system therefore settles rather than running away — the policy component drives transition, transition shrinks the policy component, and the brown share comes to rest where the forces balance.

The components of the greenium and the two investments thus form a tightly coupled system. The policy component drives both instruments — mitigation directly and transition through the valuation gap — but mitigation is misaligned, strongest where the premium is mild and absent where it is largest. The diversification component drives transition alone, because diffusion risk cannot be mitigated, only escaped. And the two investments pull the brown share in opposite directions while crowding each other out. This interplay between the priced components of the greenium and the firm's green innovation is the central mechanism of the model.

6 Conclusion

This paper has developed a unified general-equilibrium model to ask what drives the greenium and whether it spurs the green innovation it is meant to encourage. By embedding climate-policy risk, a portfolio-diversification motive, and a green-consumption preference in a single two-sector economy, the model decomposes the brown-minus-green expected-return spread into a policy, a diversification, and a hedging component. Policy risk and diversification are the primary drivers and widen the spread as the brown sector dominates, while the hedging component—rooted in the household's fear of green scarcity—runs the other way and conceals part of the spread. The

calibrated greenium lies in the range of recent robust estimates and deepens over time as climate risk intensifies, and the decomposition makes clear that the dispersion of policy shocks, not merely their average size, governs the policy component.

Endogenizing brown firms' green innovation along two margins—mitigation, which lowers the frequency of policy shocks, and transition, which converts brown capital into green—yields the paper's central result: the greenium and mitigation are misaligned. Mitigation is largest where the policy premium is mild and switches off precisely where it is largest, because once the brown sector dominates a policy shock becomes an aggregate disaster whose frequency the firm gains nothing by reducing. Operating through the stochastic discount factor, this general-equilibrium effect offers a novel, risk-based explanation for the puzzling decline in green innovation since the Paris Agreement, and shows that stricter policy can be self-defeating at the margin even as it raises brown firms' cost of capital.

The two channels and the three components of the greenium form a tightly coupled system. Because mitigation acts only on policy risk, whereas the diffusion risk behind the diversification component can be shed only by transition, the two priced risks map asymmetrically onto the two instruments. Transition tracks the valuation gap between green and brown equity until the green knowledge base that supports it thins out, and the two investments push the brown share in opposite directions while substituting for one another, so the economy settles rather than running away.

These findings carry a clear message for policy. Because the cost-of-capital channel can dull firms' incentive to abate in exactly the states where policy risk is most acute, regulation that widens the greenium need not accelerate the transition—an unintended consequence that the design of climate policy should anticipate. Given the growing importance of green consumption for the transition to sustainability, a natural next step is to let the consumption channel shape firms' innovation decisions directly, and to revisit the greenium in a setting where investors differ in their exposure to, and their preferences over, climate risk.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The Environment and Directed Technical Change," *American Economic Review*, February 2012, *102* (1), 131â66.

Aghion, Philippe, Roland Benabou, Ralf Martin, and Alexandra Roulet, "Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?," *American Economic Review: Insights*, March 2023, *5* (1), 1–20.

Ardia, David, Keven Bluteau, Kris Boudt, and Koen Inghelbrecht, "Climate Change Concerns and the Performance of Green vs. Brown Stocks," *Management Science*, 2023, *69* (12), 7607–7632.

Bansal, Ravi, Dana Kiku, and Marcelo Ochoa, "Price of Long-Run Temperature Shifts in Capital Markets," *NBER Working Papers*, 2016, (22529).

Bauer, Michael D., Daniel Huber, Glenn D. Rudebusch, and Ole Wilms, "Where is the Carbon Premium? Global Performance of Green and Brown Stocks," *Journal of Climate Finance*, 2022, *1*, 100006.

Berk, Jonathan B. and Jules H. van Binsbergen, "The impact of impact investing," *Journal of Financial Economics*, 2025, *164*, 103972.

Bolton, Patrick and Marcin Kacperczyk, "Do Investors Care about Carbon Risk?," *Journal of Financial Economics*, 2021, *142* (2), 517–549.

_ **and** _ , "Global Pricing of Carbon-Transition Risk," *The Journal of Finance*, 2023, *78* (6), 3677–3754.

Carmona, Rene, *Lectures on BSDEs, Stochastic Control, and Stochastic Differential Games with Financial Applications*, Philadelphia, PA, USA: SIAM-Society for Industrial and Applied Mathematics, 2016.

Claus, James and Jacob Thomas, "Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets," *The Journal of Finance*, 2001, *56* (5), 1629–1666.

Clements, Lee, Lily Dai, Jaakko Kooroshy, Alan Meng, and Beth Schuck, "Investing in the Green Economy: Growing in a Fractured Landscape," Technical Report, London Stock Exchange Group 2024.

Cochrane, John H, Francis A Longstaff, and Pedro Santa-Clara, "Two Trees," *The Review of Financial Studies*, 2008, *21* (1), 347–385.

Cohen, Lauren H., Umit G. Gurun, and Quoc H. Nguyen, "The ESG-Innovation Disconnect: Evidence from Green Patenting," *Management Science*, 2026, *Forthcoming.*

Colmer, Jonathan, Ralf Martin, Mirabelle Muuls, and Ulrich J Wagner, "Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System," *The Review of Economic Studies*, 05 2024, *92* (3), 1625–1660.

Dietz, Simon, Christian Gollier, and Louise Kessler, "The climate beta," *Journal of Environmental Economics and Management*, 2018, *87*, 258–274.

Duffie, Darrell and Costis Skiadas, "Continuous-Time Security Pricing: A Utility Gradient Approach," *Journal of Mathematical Economics*, 1994, *23* (2), 107–131.

Easton, Peter D., "PE Ratios, PEG Ratios, and Estimating the Implied Expected Rate of Return on Equity Capital," *The Accounting Review*, 2004, *79* (1), 73–95.

Eberly, Janice and Neng Wang, "Reallocating and Pricing Illiquid Capital: Two Productive Trees," *International Economic Review*, 2025, *66* (5), 1887–1907.

Edmans, Alex, Tom Gosling, and Dirk Jenter, "Sustainable Investing in Practice: Objectives, Beliefs, and Limits to Impact," *ECGI Finance Working Paper No. 1028/2024*, August 2025.

Eskildsen, Marc, Markus Ibert, Theis Ingerslev Jensen, and Lasse Heje Pedersen, "In Search of the True Greenium," *Available at SSRN 4744608*, 2024.

Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan, "Toward an Implied Cost of Capital," *Journal of Accounting Research*, 2001, *39* (1), 135–176.

Giese, Guido, Zoltan Nagy, and Bruno Rauis, "Foundations of Climate Investing: How Equity Markets Have Priced Climate-Transition Risks," *The Journal of Portfolio Management*, 2021, *47* (9), 35–53.

Global Sustainable Investment Alliance, "Global Sustainable Investment Review," Technical Report 2020.

Gupta, Deeksha and Jan Starmans, "Dynamic Green Disclosure Requirements," *Available at SSRN 4557187*, 2025.

Hambel, Christoph, Holger Kraft, and Frederick van der Ploeg, "Asset Diversification versus Climate Action," *International Economic Review*, 2024, *65* (3), 1323–1355.

Hartzmark, Samuel M. and Kelly Shue, "Counterproductive Sustainable Investing: The Impact Elasticity of Brown and Green Firms," *Available at SSRN 4359282*, 2023.

Heath, Davidson, Daniele Macciocchi, Roni Michaely, and Matthew C. Ringgenberg, "Does Socially Responsible Investing Change Firm Behavior?*," *Review of Finance*, 02 2023, *27* (6), 2057–2083.

Hong, Harrison, Neng Wang, and Jinqiang Yang, "Mitigating Disaster Risks in the Age of Climate Change," *Econometrica*, 2023, *91* (5), 1763–1802.

_ , _ , and _ , "Welfare Consequences of Sustainable Finance," *The Review of Financial Studies*, 2023, *36* (12), 4864–4918.

Hsu, Po-Hsuan, Kai Li, and Chi-Yang Tsou, "The Pollution Premium," *The Journal of Finance*, 2023, *78* (3), 1343–1392.

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

_ , "Growth and Ideas," in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 1B, Amsterdam: Elsevier, 2005, chapter 16, pp. 1063–1111.

Karolyi, George Andrew, Ying Wu, and Wei (William) Xiong, "Understanding the Global Equity Greenium," *Available at SSRN 4391189*, 2023.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, 03 2017, *132* (2), 665–712.

Kolbel, Julian F., Florian Heeb, Falko Paetzold, and Timo Busch, "Can Sustainable Investing Save the World? Reviewing the Mechanisms of Investor Impact," *Organization & Environment*, 2020, *33* (4), 554–574.

Krueger, Philipp, Zacharias Sautner, and Laura T Starks, "The Importance of Climate Risks for Institutional Investors," *The Review of Financial Studies*, 02 2020, *33* (3), 1067–1111.

Leippold, Markus and Tingyu Yu, "Firm-level green innovation beyond patents," *Review of Finance*, 10 2025, p. rfaf058.

Meier, Jean-Marie, Henri Servaes, Jiaying Wei, and Steven Chong Xiao, "Do Consumers Care About ESG? Evidence from Barcode-Level Sales Data," *Available at SSRN 4260716*, 2023.

Nguyen, Nam, Alejandro Rivera, and Harold H. Zhang, "Incentivizing Investors for a Greener Economy," *Journal of Financial and Quantitative Analysis*, 2024, pp. 1–41.

Ohlson, James and Beate Juettner-Nauroth, "Expected EPS and EPS growth as determinants of value," *Review of Accounting Studies*, 2005, *10*, 349–365.

Papageorgiou, Chris, Marianne Saam, and Patrick Schulte, "Substitution between Clean and Dirty Energy Inputs: A Macroeconomic Perspective," *The Review of Economics and Statistics*, 2017, *99* (2), 281–290.

Pastor, Lubos, Robert F Stambaugh, and Lucian A Taylor, "Sustainable investing in equilibrium," *Journal of Financial Economics*, 2021, *142* (2), 550–571.

_ , _ , and _ , "Dissecting Green Returns," *Journal of Financial Economics*, 2022, *146* (2), 403–424.

Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski, "Responsible investing: The ESG-efficient frontier," *Journal of Financial Economics*, 2021, *142* (2), 572–597.

Sautner, Zacharias, Laurence van Lent, Grigory Vilkov, and Ruishen Zhang, "Firm-Level Climate Change Exposure," *The Journal of Finance*, 2023, *78* (3), 1449–1498.

_ , _ , _ , and _ , "Pricing Climate Change Exposure," *Management Science*, 2023, *69* (12), 7540–7561.

Sauzet, Maxime and Olivier David Zerbib, "When Green Investors Are Green Consumers," *Available at SSRN 4098124*, 2024.

Yin, Cynthia, "Do Production Frictions Affect the Impact of Sustainable Investing?," *Fisher College of Business Working Paper No. 2024-025*, 2024.

Zerbib, Olivier David, "A Sustainable Capital Asset Pricing Model (S-CAPM): Evidence from Environmental Integration and Sin Stock Exclusion," *Review of Finance*, 07 2022, *26* (6), 1345–1388.

Zhang, Shaojun, "Carbon Returns across the Globe," *The Journal of Finance*, 2024, *80* (1), 615–645.

Appendix

Throughout, for two-vectors $u = (u_1, u_2)$ and $v = (v_1, v_2)$, I write $\langle u, v\rangle_\nu = u_1v_1 + u_2v_2 + \nu\,(u_1v_2 + u_2v_1)$ and $|u|_\nu^2 = \langle u, u\rangle_\nu$ for the covariance form induced by the Brownian correlation ν. I collect the diffusion loadings of η in $\sigma_\eta = (\eta(1-\eta)\sigma_b,\, -\eta(1-\eta)\sigma_g)$, so that $|\sigma_\eta|_\nu^2 = \eta^2(1-\eta)^2\left(\sigma_b^2 + \sigma_g^2 - 2\nu\sigma_b\sigma_g\right)$.

A.1. Household's problem

State dynamics. By definition $\eta = Y_b/(Y_b + Y_g)$. Applying Itô's lemma to the output processes (2.1)–(2.2), and using that transition is an idiosyncratic Poisson event that by the law of large numbers removes brown output at the deterministic rate ψ (so that the green sector receives the inflow $\psi\,\eta/(1-\eta)$ while only the aggregate policy shock causes a jump in η), the brown share evolves as

$$d\eta = \widehat{\Sigma}(\eta)\,dt + \eta(1-\eta)\left[\sigma_b dB_b - \sigma_g dB_g\right] + \left[\eta^J(Z) - \eta\right] dN_p, \tag{6.1}$$

$$\widehat{\Sigma}(\eta) = \eta(1-\eta)\Big[\mu_b - z_m - z_t - \mu_g - \eta\sigma_b^2 + (1-\eta)\sigma_g^2 + (2\eta-1)\nu\sigma_b\sigma_g\Big] - \psi\,\eta, \tag{6.2}$$

where the $-\psi\eta$ term is the net of the brown outflow $-\psi$ and the green inflow $+\psi\eta/(1-\eta)$, and $\eta^J(Z) = \eta Z/[\eta Z + (1-\eta)]$.

Valuation ratios and returns. By Itô's lemma, the brown price/output ratio satisfies

$$dp_b(\eta) = \left[p_b'(\eta)\,\widehat{\Sigma}(\eta) + \tfrac{1}{2}p_b''(\eta)\,|\sigma_\eta|_\nu^2\right] dt + p_b'(\eta)\,\eta(1-\eta)\left[\sigma_b dB_b - \sigma_g dB_g\right] + \left[p_b(\eta^J) - p_b(\eta)\right] dN_p. \tag{6.3}$$

The brown firm value is $P_b = Y_b\,p_b(\eta)$, and brown output carries the transition outflow, $dY_b/Y_b = (\mu_b - z_m - z_t - \psi)\,dt + \sigma_b dB_b + (Z-1)dN_p$. Computing $d(Y_b p_b)$, adding the dividend yield $[m_b(\eta) + \psi\,p_g(\eta)]\,/p_b(\eta)$ (output sold plus the spinoff value of newly transitioned green equity), and collecting the deterministic transition pieces $\psi\,p_g/p_b - \psi$ into $\psi\,[p_g - p_b]^+/p_b$, the brown return is

$$dR_b = \Lambda_b(\eta)\,dt + \Gamma_b(\eta)\,dB_b - \Delta_b(\eta)\,dB_g - \left[1 - \frac{Z\,p_b(\eta^J)}{p_b(\eta)}\right] dN_p, \tag{6.4}$$

where

$$\Lambda_b(\eta) = \frac{m_b(\eta)}{p_b(\eta)} + (\mu_b - z_m - z_t) + \psi\,\frac{[p_g(\eta) - p_b(\eta)]^+}{p_b(\eta)} + \frac{p_b'(\eta)}{p_b(\eta)}\widehat{\Sigma}(\eta) + \tfrac{1}{2}\frac{p_b''(\eta)}{p_b(\eta)}\,|\sigma_\eta|_\nu^2 + \frac{p_b'(\eta)}{p_b(\eta)}\eta(1-\eta)\sigma_b\,(\sigma_b - \nu\sigma_g)\,. \tag{6.5}$$

The green firm's existing equity grows organically at μ_g (the transition inflow accrues as newly created equity to brown shareholders, not to incumbent green holders, and so enters the green return only through the state η in $\widehat{\Sigma}$), and green output is not directly struck by the policy shock. Following the same steps,

$$dR_g = \Lambda_g(\eta)\,dt + \Gamma_g(\eta)\,dB_b + \Delta_g(\eta)\,dB_g - \left[1 - \frac{p_g(\eta^J)}{p_g(\eta)}\right] dN_p, \tag{6.6}$$

$$\Lambda_g(\eta) = \frac{m_g(\eta)}{p_g(\eta)} + \mu_g + \frac{p_g'(\eta)}{p_g(\eta)}\widehat{\Sigma}(\eta) + \tfrac{1}{2}\frac{p_g''(\eta)}{p_g(\eta)}\,|\sigma_\eta|_\nu^2 + \frac{p_g'(\eta)}{p_g(\eta)}\eta(1-\eta)\sigma_g\,(\nu\sigma_b - \sigma_g)\,, \tag{6.7}$$

where the diffusion loadings are

$$\Gamma_b = \sigma_b\left[1 + \frac{p_b'}{p_b}\eta(1-\eta)\right], \quad \Delta_b = \sigma_g\frac{p_b'}{p_b}\eta(1-\eta), \quad \Gamma_g = \sigma_b\frac{p_g'}{p_g}\eta(1-\eta), \quad \Delta_g = \sigma_g\left[1 - \frac{p_g'}{p_g}\eta(1-\eta)\right], \tag{6.8}$$

so that $\sigma_{R_b} = (\Gamma_b, -\Delta_b)$ and $\sigma_{R_g} = (\Gamma_g, \Delta_g)$.

Wealth dynamics and the HJB. With $c_n = C_n/W$ and $\sigma_W = \pi_b\sigma_{R_b} + \pi_g\sigma_{R_g}$, the budget constraint is

$$\frac{dW}{W} = [r - c_b m_b - c_g m_g + \pi_b\,(\Lambda_b - r) + \pi_g\,(\Lambda_g - r)]\,dt + \sigma_{W,b}\,dB_b + \sigma_{W,g}\,dB_g + J_W\,dN_p, \tag{6.9}$$

with $J_W = \pi_b\big[Z\,p_b(\eta^J)/p_b - 1\big] + \pi_g\big[p_g(\eta^J)/p_g - 1\big]$. With Duffie–Epstein–Zin utility under unit EIS, $f(C,V) = \rho(1-\gamma)V\Big[\log C - \frac{1}{1-\gamma}\log((1-\gamma)V)\Big]$, and the household's HJB is

$$0 = \sup_{C_b, C_g, \pi_b, \pi_g}\Big\{f(C,V) + V_W W[r - c_b m_b - c_g m_g + \pi_b(\Lambda_b - r) + \pi_g(\Lambda_g - r)] + \tfrac{1}{2}V_{WW}W^2\,|\sigma_W|_\nu^2 + V_\eta\widehat{\Sigma} + \tfrac{1}{2}V_{\eta\eta}\,|\sigma_\eta|_\nu^2 + V_{W\eta}W\,\langle\sigma_W, \sigma_\eta\rangle_\nu + \widetilde{\lambda}\,\mathbb{E}_Z\big[V^J - V\big]\Big\}, \tag{6.10}$$

where $\widetilde{\lambda} = (\lambda_g + \eta\Delta_\lambda)/(1+\phi z_m)$, $V^J = V(W^J;\eta^J)$, and $W^J = W\,(1+J_W)$; the idiosyncratic transition does not appear because it is diversifiable. I guess and verify

$$V(W;\eta) = \rho^{1-\gamma} G(\eta) \frac{W^{1-\gamma}}{1-\gamma}. \tag{6.11}$$

Consumption. Substituting (6.11) into (6.10), the first-order conditions for C_b, C_g are

$$\frac{\rho W\,(1-\theta)^{\frac{1}{\alpha}} C_b^{-\frac{1}{\alpha}}}{(1-\theta)^{\frac{1}{\alpha}} C_b^{\frac{\alpha-1}{\alpha}} + \theta^{\frac{1}{\alpha}} C_g^{\frac{\alpha-1}{\alpha}}} = m_b, \qquad \frac{\rho W\,\theta^{\frac{1}{\alpha}} C_g^{-\frac{1}{\alpha}}}{(1-\theta)^{\frac{1}{\alpha}} C_b^{\frac{\alpha-1}{\alpha}} + \theta^{\frac{1}{\alpha}} C_g^{\frac{\alpha-1}{\alpha}}} = m_g, \tag{6.12}$$

which solve to

$$C_b = \frac{\rho W}{m_b\left[1 + \xi^{1-\alpha}\left(\frac{\theta}{1-\theta}\right)\right]}, \qquad C_g = \xi^{-\alpha}\left(\tfrac{\theta}{1-\theta}\right) C_b. \tag{6.13}$$

It follows that the aggregate consumption claim has a constant price/dividend ratio, $m_b C_b + m_g C_g = \rho W$, and by the basket normalization $C = \rho W$.

Portfolio. Substituting (6.13) and (6.11) into (6.10) reduces it to the equation for G,

$$\begin{aligned} 0 = \sup_{\pi_b,\pi_g} \Big\{ &-\rho \log G + (1-\gamma)\left[r - \rho + \pi_b(\Lambda_b - r) + \pi_g(\Lambda_g - r)\right] - \tfrac{1}{2}\gamma(1-\gamma)\,|\sigma_W|_\nu^2 \\ &+ (1-\gamma)\frac{G'}{G}\langle \sigma_W, \sigma_\eta\rangle_\nu + \frac{G'}{G}\widehat{\Sigma} + \tfrac{1}{2}\frac{G''}{G}\,|\sigma_\eta|_\nu^2 + \widetilde{\lambda}\,\mathbb{E}_Z\left[\left(\tfrac{W^J}{W}\right)^{1-\gamma}\frac{G(\eta^J)}{G(\eta)} - 1\right]\Big\}. \end{aligned} \tag{6.14}$$

The first-order conditions for π_b, π_g give the optimal portfolio

$$\Sigma(\eta)\begin{bmatrix}\pi_b \\ \pi_g\end{bmatrix} = \frac{1}{\gamma}\left\{\begin{bmatrix}\Lambda_b - r \\ \Lambda_g - r\end{bmatrix} + \widetilde{\lambda}\begin{bmatrix}\mathbb{E}_Z\left[\kappa_p\left(Z\,p_b(\eta^J)/p_b - 1\right)\right] \\ \mathbb{E}_Z\left[\kappa_p\left(p_g(\eta^J)/p_g - 1\right)\right]\end{bmatrix} + \Omega(\eta)\frac{G'(\eta)}{G(\eta)}\right\}, \tag{6.15}$$

where κ_p is the price of policy risk derived in A.3,

$$\Sigma(\eta) = \begin{bmatrix} |\sigma_{R_b}|_\nu^2 & \langle \sigma_{R_b}, \sigma_{R_g}\rangle_\nu \\ \langle \sigma_{R_b}, \sigma_{R_g}\rangle_\nu & |\sigma_{R_g}|_\nu^2 \end{bmatrix}, \qquad \Omega(\eta) = \begin{bmatrix}\langle \sigma_{R_b}, \sigma_\eta\rangle_\nu \\ \langle \sigma_{R_g}, \sigma_\eta\rangle_\nu\end{bmatrix}. \tag{6.16}$$

Written out, Σ and Ω are

$$\Sigma(\eta) = \begin{bmatrix} \Gamma_b^2 + \Delta_b^2 - 2\nu\Gamma_b\Delta_b & \Gamma_b\Gamma_g - \Delta_b\Delta_g + \nu(\Gamma_b\Delta_g - \Delta_b\Gamma_g) \\ \Gamma_b\Gamma_g - \Delta_b\Delta_g + \nu(\Gamma_b\Delta_g - \Delta_b\Gamma_g) & \Gamma_g^2 + \Delta_g^2 + 2\nu\Gamma_g\Delta_g \end{bmatrix},$$

$$\Omega(\eta) = \eta(1-\eta) \begin{bmatrix} (\sigma_b - \nu\sigma_g)\Gamma_b + (\sigma_g - \nu\sigma_b)\Delta_b \\ (\sigma_b - \nu\sigma_g)\Gamma_g - (\sigma_g - \nu\sigma_b)\Delta_g \end{bmatrix}.$$

A.2. Brown firm's problem

The brown firm takes the SDF dynamics (3.14) as given (its prices of risk Φ_b, Φ_g, κ_p are derived in A.3). No arbitrage requires that $[m_b + \psi p_g]\, M_t Y_{b,t}\, dt + d\,(M_t P_{b,t})$ be a martingale, i.e. have zero drift under the physical measure. Setting that drift to zero yields the firm's HJB,

$$\begin{aligned} r\, p_b = \sup_{z_m, z_t \geq 0} \; & m_b + p_b\,(\mu_b - z_m - z_t) + \psi(z_t)\,[p_g - p_b]^+ + p_b' \left[\widehat{\Sigma} + \eta(1-\eta)\sigma_b(\sigma_b - \nu\sigma_g)\right] + \tfrac{1}{2} p_b'' \,|\sigma_\eta|_\nu^2 \\ & - p_b\,\langle \Phi, \sigma_{R_b} \rangle_\nu - p_b\,\widetilde{\lambda}\,\mathbb{E}_Z[\kappa_p - 1] + \widetilde{\lambda}\,\mathbb{E}_Z\left[\kappa_p\, Z\, p_b(\eta^J) - p_b\right], \end{aligned} \tag{6.17}$$

where $\Phi = (\Phi_b, \Phi_g)$ and $\widetilde{\lambda} = (\lambda_g + \eta\Delta_\lambda)/(1 + \phi z_m)$. The firm internalizes its effect on its own jump term $\widetilde{\lambda}\,\mathbb{E}_Z[\kappa_p Z p_b^J - p_b]$ but takes the equilibrium SDF compensator $-p_b\widetilde{\lambda}\,\mathbb{E}_Z[\kappa_p - 1]$ as given. Retaining only the z_m, z_t-dependent terms, the firm solves

$$\{z_m, z_t\} = \arg\max_{z_m, z_t \geq 0} \; p_b\,(1 - z_m - z_t) + \left(\frac{\lambda_g + \eta\Delta_\lambda}{1 + \phi z_m}\right) \mathbb{E}_Z\left[\kappa_p\, Z\, p_b(\eta^J) - p_b\right] + \psi(z_t)\,[p_g - p_b]^+ . \tag{6.18}$$

The first-order condition in z_m, using $\partial\widetilde{\lambda}/\partial z_m = -\phi(\lambda_g + \eta\Delta_\lambda)/(1 + \phi z_m)^2$, gives

$$z_m = \max\left\{0,\; \frac{1}{\phi}\left(\sqrt{\frac{\phi(\lambda_g + \eta\Delta_\lambda)}{p_b}\Big[\, p_b - \mathbb{E}_Z(\kappa_p\, Z\, p_b(\eta^J))\,\Big]^+} - 1\right)\right\}, \tag{6.19}$$

and the first-order condition in z_t, using $\psi(z_t) = \bar{\psi}\, z_t^{\beta_\psi}(1-\eta)^\kappa$ so that $\psi'(z_t) = \bar{\psi}\,\beta_\psi z_t^{\beta_\psi - 1}(1-\eta)^\kappa$, gives

$$z_t = \left[\bar{\psi}\,\beta_\psi (1-\eta)^\kappa\, \frac{[p_g - p_b]^+}{p_b}\right]^{\frac{1}{1-\beta_\psi}}. \tag{6.20}$$

A.3. Competitive market equilibrium

The SDF. Since $f_C(C,V) = \rho(1-\gamma)V/C = \rho^{1-\gamma}G(\eta)W^{-\gamma}$ and $f_V(C,V) = -\rho\,[1+\log G(\eta)] \equiv h(\eta)$, Itô's lemma applied to $M_t = e^{\int_0^t f_V\,ds} f_C$ gives $dM/M = h(\eta)\,dt + df_C/f_C$. Expanding df_C/f_C and collecting terms,

$$\frac{dM}{M^-} = -r(\eta)\,dt - \Phi_b(\eta)\,dB_b - \Phi_g(\eta)\,dB_g + (\kappa_p - 1)\left[dN_p - \widetilde{\lambda}(\eta)\,dt\right], \tag{6.21}$$

with diffusion prices of risk

$$\Phi_b = \gamma\,\sigma_{W,b} - \frac{G'}{G}\,\sigma_{\eta,b}, \qquad \Phi_g = \gamma\,\sigma_{W,g} - \frac{G'}{G}\,\sigma_{\eta,g}, \tag{6.22}$$

and the price of policy risk

$$\kappa_p(\eta, Z) = \left(\frac{W^J}{W}\right)^{-\gamma}\frac{G(\eta^J)}{G(\eta)}, \qquad \frac{W^J}{W} = \pi_b\frac{Z\,p_b(\eta^J)}{p_b} + \pi_g\frac{p_g(\eta^J)}{p_g}. \tag{6.23}$$

The idiosyncratic transition is diversifiable and so does not enter (6.21).

Risk-free rate. Matching the drift of (6.21) against $h(\eta)\,dt + df_C/f_C$ yields the raw expression

$$\begin{aligned} r(\eta) = {} & \rho\,[1-\gamma+\log G] + \gamma\,[\pi_b\Lambda_b + \pi_g\Lambda_g] - \frac{G'}{G}\widehat{\Sigma} - \tfrac{1}{2}\frac{G''}{G}\,|\sigma_\eta|_\nu^2 \\ & - \tfrac{1}{2}\gamma(\gamma+1)\,|\sigma_W|_\nu^2 + \gamma\frac{G'}{G}\,\langle\sigma_W, \sigma_\eta\rangle_\nu - \widetilde{\lambda}\,\mathbb{E}_Z[\kappa_p - 1]\,. \end{aligned} \tag{6.24}$$

Adding this to the reduced-form equation (6.14) (which holds with equality at the optimal portfolio) cancels the G-derivative terms and delivers the familiar decomposition,

$$r(\eta) = \underbrace{\rho + \mu_c(\eta) - \gamma\sigma_c^2(\eta)}_{\text{Diffusion Risks}} + \underbrace{\widetilde{\lambda}(\eta)\,\mathbb{E}_Z\left[\kappa_p\left(\pi_b\frac{Z\,p_b(\eta^J)}{p_b} + \pi_g\frac{p_g(\eta^J)}{p_g} - 1\right)\right]}_{\text{Policy Risk}}, \tag{6.25}$$

where $\mu_c(\eta) = \pi_b\Lambda_b + \pi_g\Lambda_g - \rho$ and $\sigma_c^2(\eta) = |\sigma_W|_\nu^2$ are the drift and variance of consumption growth.

Relative price and the pricing system. Market clearing $C_b = Y_b$, $C_g = Y_g$ combined with (6.13) gives $C_b/C_g = \xi^\alpha \left(\frac{1-\theta}{\theta}\right) = \eta/(1-\eta)$, i.e.

$$\xi(\eta) = \left(\frac{\theta}{1-\theta}\right)^{\frac{1}{\alpha}} \left(\frac{\eta}{1-\eta}\right)^{\frac{1}{\alpha}}. \tag{6.26}$$

The aggregate-consumption claim condition $[Y_b p_b + Y_g p_g] / [Y_b m_b + Y_g m_g] = 1/\rho$ pins p_g residually:

$$p_g(\eta) = \frac{1}{\rho}\left[\frac{\eta}{1-\eta} m_b(\eta) + m_g(\eta)\right] - \frac{\eta}{1-\eta} p_b(\eta). \tag{6.27}$$

Substituting (6.27) into the brown HJB (6.17) and the G-equation (6.14) yields a system of delay ODEs in $p_b(\eta)$ and $G(\eta)$, solved numerically under the boundary conditions

$$G(0) = \exp\left\{\frac{(1-\gamma)\left(2\mu_g - \gamma\sigma_g^2\right)}{2\rho}\right\}, \tag{6.28}$$

$$G(1) = \exp\left\{\frac{2\lambda_b\left(\mathbb{E}_Z\left[Z^{1-\gamma}\right] - 1\right) + (1-\gamma)\left(2\mu_b - \gamma\sigma_b^2\right)}{2\rho}\right\}, \tag{6.29}$$

$$p_b(0) = +\infty, \qquad p_b(1) = \frac{(1-\theta)^{\frac{1}{\alpha-1}}}{\rho}, \tag{6.30}$$

where $\lambda_b = \lambda_g + \Delta_\lambda$ is the policy intensity in the fully brown state.

A.4. Asset prices

By the fundamental theorem of asset pricing,

$$\frac{\mathbb{E}[dR_i]}{dt} - r = -\frac{\mathbb{E}\left[dR_i \frac{dM}{M}\right]}{dt}, \qquad i \in \{b, g\}. \tag{6.31}$$

From (6.21), the SDF loads $-\Phi_b, -\Phi_g$ on the Brownian shocks and jumps by $\kappa_p - 1$. The diffusion part of the covariance contributes $\langle \sigma_{R_i}, \Phi \rangle_\nu$, while the compensated jump contributes $-\widetilde{\lambda}\,\mathbb{E}_Z[(\kappa_p - 1)(\mathcal{J}_i - 1)]$, where $\mathcal{J}_b = Z\, p_b(\eta^J)/p_b$ and $\mathcal{J}_g = p_g(\eta^J)/p_g$ are the multiplicative jumps in the two

returns. Hence the equity premia are

$$\frac{\mathbb{E}[dR_b]}{dt} - r = \langle \sigma_{R_b}, \Phi \rangle_\nu - \widetilde{\lambda}\,\mathbb{E}_Z\big[(\kappa_p - 1)(\mathcal{J}_b - 1)\big], \tag{6.32}$$

$$\frac{\mathbb{E}[dR_g]}{dt} - r = \big\langle \sigma_{R_g}, \Phi \big\rangle_\nu - \widetilde{\lambda}\,\mathbb{E}_Z\big[(\kappa_p - 1)(\mathcal{J}_g - 1)\big]. \tag{6.33}$$

Substituting the prices of risk $\Phi = \gamma\sigma_W - (G'/G)\sigma_\eta$ splits each diffusion premium into a market term $\gamma\,\langle \sigma_{R_i}, \sigma_W \rangle_\nu$ and a hedging term $-(G'/G)\,\langle \sigma_{R_i}, \sigma_\eta \rangle_\nu$, which are illustrated in Figure 9. Collecting the covariance loadings of each return,

$$\widehat{\Theta}_b = \begin{bmatrix} \Gamma_b - \nu\Delta_b \\ \nu\Gamma_b - \Delta_b \end{bmatrix}, \quad \widehat{\Theta}_g = \begin{bmatrix} \Gamma_g + \nu\Delta_g \\ \nu\Gamma_g + \Delta_g \end{bmatrix}, \quad \Theta_W = \begin{bmatrix} \pi_b\Gamma_b + \pi_g\Gamma_g \\ -\pi_b\Delta_b + \pi_g\Delta_g \end{bmatrix}, \quad \Theta_G = \begin{bmatrix} -\frac{G'}{G}\eta(1-\eta)\sigma_b \\ \frac{G'}{G}\eta(1-\eta)\sigma_g \end{bmatrix},$$

the diffusion premium of claim i is $\gamma\,\widehat{\Theta}_i^\top\Theta_W + \widehat{\Theta}_i^\top\Theta_G$, with $\widehat{\Theta}_b, \widehat{\Theta}_g$ the quantities of diffusion risk shown in Figure 12.

Differencing the two premia gives the greenium. The shared market and hedging structure collapses onto the return spread $\sigma_{R_b} - \sigma_{R_g} = (1 + h_b - h_g)(\sigma_b, -\sigma_g)$, where $h_i = (p_i'/p_i)\eta(1-\eta)$, and $\langle (\sigma_b, -\sigma_g), \sigma_\eta \rangle_\nu = \eta(1-\eta)(\sigma_b^2 + \sigma_g^2 - 2\nu\sigma_b\sigma_g)$. Substituting Φ yields the three-component decomposition (4.4): the diversification component $\gamma\,\big\langle \sigma_{R_b} - \sigma_{R_g}, \sigma_W \big\rangle_\nu$, the hedging component $-(G'/G)\,\big\langle \sigma_{R_b} - \sigma_{R_g}, \sigma_\eta \big\rangle_\nu$, and the policy component $-\widetilde{\lambda}\,\mathbb{E}_Z[(\kappa_p - 1)(\mathcal{J}_b - \mathcal{J}_g)]$.