Betting on Green: A Green Security Market Line

Zeynep Gurur Akbas* Enrico De Giorgi†

February 1, 2026

Abstract

Sustainable investing is increasingly recognized as a market-based mechanism for reallocating capital toward greener firms by influencing their cost of capital. However, the effect between a firm's greenness and its discount rates is unclear. Here, we develop a Green Security Market Line that separates greenness channels by augmenting standard cross-sectional pricing tests with (i) an ESG-level component and (ii) a firm-level green beta that captures how strongly an ESG profile moves with aggregate market greenness. Using U.S. equities from 2007–2023 and MSCI ESG ratings, we find that higher ESG levels and higher green betas are associated with lower required returns. The exposure effect is driven largely by within-industry dispersion rather than industry tilts, whereas time-variation tests reveal a *rotation* in sustainability pricing. When climate attention spikes, the discount associated with green-beta exposure compresses, while the valuation wedge tied to ESG levels widens. The results clarify that greenium is not a single object, but the net effect of a level (taste/characteristic) component and a systematic exposure component.

JEL: G12, G14, Q54.

Keywords: Green beta; Greenium; Climate risk; ESG; Implied cost of capital.

*Department of Economics, School of Economics and Political Science, University of St. Gallen, Bodanstrasse 6, 9000 St. Gallen, Switzerland; (corresponding author) zeynepgurur.akbas@unisg.ch

†Department of Economics, School of Economics and Political Science, University of St. Gallen, Bodanstrasse 6, 9000 St. Gallen, Switzerland; ORCID: 0000-0001-8017-9618; enrico.degiorgi@unisg.ch

1 Introduction

Sustainability has moved from the periphery of finance to a first-order force shaping capital allocation. Institutional mandates, regulation, and investor demand increasingly link firms' financing conditions to environmental, social, and governance (ESG) characteristics. Yet, a central asset-pricing question remains unsettled: do "green" firms face lower required returns because investors *like* holding them or because they provide a *hedge* against aggregate sustainability shocks and climate risks? We use the term *greenium* to denote the expected return differential associated with greenness; a negative greenium means that greener firms command lower expected returns (and a lower cost of equity) than otherwise comparable firms. Empirically, answering this question requires separating pricing of sustainability *levels* from pricing of sustainability *exposures*.

Disentangling these channels empirically presents a significant identification challenge, as ESG levels and systematic exposures exhibit high collinearity in historical data. Evidence based on realized returns is noisy at monthly horizons, and ESG ratings differ across providers and embed substantial disagreement (Berg et al., 2022; Christensen et al., 2022; Gibson et al., 2021). Following Eskildsen et al. (2024), we select the implied cost of capital (ICC), a forward-looking proxy for expected returns based on prices and cash-flow expectations (Claus and Thomas, 2001; Gebhardt et al., 2001; Easton, 2004; Ohlson and Jüttner-Nauroth, 2005), as our main pricing target. We combine ICC with a new firm-level measure of systematic sustainability exposure that we call *green beta*. The green beta, $\beta^g_{i,t}$, is defined as the sensitivity of a firm's ESG score to market-wide movements in ESG, estimated from a regression in rolling windows (see (5.2)).

Intuitively, $\beta^g_{i,t}$ captures whether a firm's sustainability profile moves with aggregate market greenness. That is, firms with high $\beta^g_{i,t}$ "become greener" when the market becomes greener, while low-$\beta^g_{i,t}$ firms move less with aggregate greenness. This exposure concept is distinct from a firm's ESG *level* $s_{i,t}$: two firms may have similar ESG levels

yet differ substantially in how their ESG profiles covary with aggregate sustainability dynamics. The distinction is important since investors can value static greenness (a taste channel) while also valuing exposure that provides diversification or hedging benefits (an exposure channel).

We formalize this concept in an equilibrium model that delivers a *Green Security Market Line* (GSML). The GSML extends the Capital Asset Pricing Model (CAPM) by allowing investors to care directly about greenness and by introducing a second beta, β^g, alongside the market beta β^m. Expected returns can fall with ESG levels if some investors accept lower returns to hold greener firms, and can also fall with β^g if exposure to aggregate greenness is valued as a hedge or diversification attribute. This structure delivers a two-channel interpretation: ESG levels behave like a characteristic, while green beta captures exposure to an aggregate sustainability state variable that can be priced. Under *pure taste*, only levels should be priced; under *pure hedging/risk*, only exposures should be priced. Our evidence speaks to which forces are present and how their strength varies over time.

We study U.S. equities from 2007–2023, combining monthly MSCI ESG scores with implied costs of capital (ICC) computed from analysts' forecasts. We estimate monthly Fama–MacBeth cross-sections with standard controls (size, value, and past returns) and complement the regressions with portfolio sorts, factor-mimicking portfolios, and spanning tests using Fama–French factors. To probe state dependence, we interact β^g and ESG levels with the Media Climate Change Concern (MCCC) index (Ardia et al., 2021), in the spirit of climate-salience tests in Pástor et al. (2022). Finally, to address the concern that green-beta pricing might simply reflect persistent sector composition (green industries versus brown industries), we conduct industry-portfolio tests, a within–between decomposition, and within-industry neutral portfolio sorts that hold industry weights fixed by construction.

The results support three main conclusions. First, green beta is priced in the cross-section of ICCs: firms with higher β^g have lower implied costs of capital, even

after controlling for ESG levels and standard characteristics. Second, ESG levels are strongly priced: higher-$s_{i,t}$ firms exhibit materially lower ICCs, consistent with a sizeable preference-based discount. Third, climate salience does not simply scale the greenium up or down; it *reallocates* it across channels. In high-attention months, the exposure-related component is less negative, while the level-related discount deepens. Section 6.2 quantifies this change using monthly prices extracted from the ICC cross-sections and innovations in the MCCC climate-concern index. Moreover, this reweighting is not confined to high-frequency shocks. Consistent with the growing salience of sustainability rankings and institutional climate-risk management, the estimated exposure price $\lambda_{g,t}$ trends upward (toward zero) after the mid-2010s, even as the characteristic-based component steepens. This pattern aligns with evidence that the publication of salient sustainability ratings triggered large reallocations of fund flows (Hartzmark and Sussman, 2019) and with survey evidence that large institutional investors increasingly incorporate climate risks into portfolio and risk-management decisions (Krueger et al., 2020).

A key contribution of the paper is to show that the green-beta channel is not an industry-tilt story. Industry portfolios are a coarse aggregation that mechanically averages away within-industry dispersion in β^g; accordingly, the incremental industry-level green-beta slope attenuates once we saturate the industry cross-section with standard characteristics. However, when we decompose $\beta^g_{i,t}$ into within- and between-industry components and construct within-industry neutral portfolios, the negative pricing of β^g remains concentrated within industries. This evidence reconciles the strong firm-level effect with weaker incremental industry-portfolio slopes and strengthens the interpretation of green beta as a diversification-relevant attribute that varies *within* sectors.

Ultimately, our contribution to the green asset-pricing literature is to (i) provide a GSML that separates exposure from level channels, (ii) introduce a measurement strategy that combines a firm-level green beta with forward-looking expected-return proxies

(ICC), and (iii) document cross-sectional, portfolio, and industry-decomposition evidence that both channels matter and respond differently to climate-salience variation. The remainder of the paper proceeds as follows: Section 2 reviews related literature, Section 3 develops the model, Sections 4–5 describe data and methodology, Section 6 presents results, and Section 7 concludes.

2 Literature Review

Sustainable investing is now large enough to plausibly affect equilibrium discount rates, but the sign of "greenium" remains contested. Preference-based models imply that investors bid up greener firms and accept lower expected returns, while risk-based views emphasise compensation for exposure to transition, liability, and regulatory risks borne disproportionately by carbon-intensive firms. Empirically, distinguishing these channels requires separating the pricing of ESG levels from the pricing of ESG exposures, while accounting for rating disagreement, capital flows, and trading frictions that can obscure expected-return signals. For a recent synthesis of mechanisms and evidence in sustainable investing (including tastes, tilts, and greenium), see Pástor et al. (2025).

Two benchmark equilibrium frameworks clarify how sustainability can enter prices. Pedersen et al. (2021) model ESG as both information about fundamentals and a direct argument in investors' utility, generating an ESG-efficient frontier and an ESG-adjusted CAPM. Pástor et al. (2021) also embed non-pecuniary ESG preferences in general equilibrium, implying lower expected returns for greener assets; in addition, shocks to an "ESG taste" component, which captures shifts in investor appetite for sustainability, can generate temporary green outperformance. Together, these models highlight a natural distinction between a level channel (valuing high ESG characteristics) and an exposure channel (valuing assets that covary with aggregate sustainability states).

Subsequent extensions introduce heterogeneous preferences, rating uncertainty,

and market segmentation. Avramov et al. (2022) introduce "ESG rating uncertainty" as a priced risk factor: greater dispersion across rating agencies raises the aggregate market risk premium and a stock's effective beta, thereby weakening the negative relation between ESG scores and alpha. Zerbib (2022) develops a "sustainable CAPM" that separates an exclusion premium from a taste premium that compensates investors holding lower-ESG portfolios. Hens and Trutwin (2024) offer a parsimonious extension in which investors demand higher expected returns for poor-ESG stocks and wealthier investors place greater weight on ESG criteria. Additionally, Johnstone and Grant (2023) show that even an ESG project with zero beta and a non-positive payoff can increase the firm's required return. Dobrick et al. (2025) treat ESG as a factor by constructing zero-investment portfolios based on multiple ESG data providers; adding these ESG factors improves the explanatory power of expanded Fama–French models in both North America and Europe. Similarly, Nakagawa et al. (2025) incorporate stochastic ESG scores and non-pecuniary utility into a two-beta model, showing that expected returns depend on both market and ESG betas; their framework allows CFOs to estimate a firm's cost of capital using these two betas. Finally, Hara and Hens (2025) analyze shareholder engagement in an ESG-CAPM with incomplete markets, finding that its impact on equilibrium prices is negligible. These models imply that the sign of the greenium depends on the balance between (i) preference-driven demand and (ii) risk compensation for climate-related shocks.

Empirical research tests these theoretical predictions using two-pass Fama–MacBeth regressions. The typical approach first estimates each stock's exposure (beta) to a green factor, then regresses realized or expected returns on those green betas to infer the price of green risk. Under the risk-compensation hypothesis, brown firms (with low ESG) earn higher expected returns to compensate for risk. Under the taste-based view, green firms offer non-pecuniary benefits to some investors, who are willing to accept lower financial returns. In Pástor et al. (2022), they construct a high-minus-low (green-minus-brown) portfolio using MSCI environmental ratings. They show that controlling

for this green factor dramatically reduces the abnormal returns (alphas) of traditional Fama–French value and momentum portfolios. To interpret this finding, the authors develop a monthly climate-concern index based on the frequency of climate-related terms in news articles. They demonstrate that green factor returns are positively related to unexpected increases in climate concern, suggesting that realized green outperformance has been driven by shocks to investor tastes. In their framework, the expected price of green exposure is negative and investors accept lower expected returns to hold assets that hedge climate-related states. When the green factor is included, the previously observed value and momentum anomalies largely disappear. Their evidence suggests that climate-related state variables help explain variation in green-versus-brown performance that standard factor models do not capture.

Using a broader measurement set, Eskildsen et al. (2024) evaluate 23 greenness proxies across global equity markets. Across many portfolio implementations, most realised green-minus-brown estimates are close to zero once multiple-testing concerns are addressed. Aggregating signals across measures reduces noise, and using implied costs of capital as an expected-return proxy yields a modest but economically meaningful negative price of green exposure (on the order of tens of basis points per standard deviation), with larger discounts in markets with stronger sustainability preferences.

The stability of ESG premia is complicated by divergence across rating agencies. Berg et al. (2022) compare ESG ratings from six major providers and find that pairwise correlations range only from 0.38 to 0.71. They further decompose this divergence into differences in measurement (accounting for about 56% of the variation), scope (38%), and weighting (6%), indicating that agencies focus on different indicators and methodologies. Christensen et al. (2022) show that when firms increase their ESG disclosure, rating disagreement actually grows, specifically for outcome-based metrics which is associated with higher stock return volatility and a lower likelihood of external financing. Importantly, rating disagreement itself appears to be priced in the market. Gibson et al. (2021) document that stock returns are positively related to ESG rating

dispersion, implying a risk premium for firms with more uncertain ESG assessments. Consistent with this, the model of Avramov et al. (2022) predicts that higher rating uncertainty raises the overall market risk premium and increases a stock's alpha. These findings underscore the value of using aggregated scores or factor-mimicking portfolios to reduce measurement noise. Indeed, the composite green score developed by Eskildsen et al. (2024) helps to stabilize signals and materially improves inference.

Risk-compensation theories receive support from studies focusing on environmental exposures. Bolton and Kacperczyk (2021) find that firms with higher carbon dioxide emissions earn higher subsequent returns even after controlling for standard risk factors and profitability, suggesting that investors demand extra compensation for carbon-related liabilities. In a related study, Oestreich and Tsiakas (2015) examine the first phase of the European Union's Emissions Trading Scheme (EU ETS) and find that firms receiving free carbon allowance allocations significantly outperformed those without such allowances; accordingly, a carbon factor helps to explain cross-sectional return differences in that setting. Li et al. (2023) document a sizable "pollution premium": a long–short portfolio based on toxic emission intensity delivers an average annual return of roughly 4.4% that is not explained by standard risk factors. They propose a new systematic risk factor linked to environmental policy uncertainty. Collectively, these studies imply that carbon emissions and other pollution exposures can command positive return premia, especially in contexts where liability risks and regulatory pressures are salient.

Because climate risk is dynamic, static emissions metrics may be insufficient to capture it. Engle et al. (2020) construct a climate news index by parsing newspaper articles for climate-related terms and extracting innovations in climate news. Using a mimicking-portfolio approach, they build hedge portfolios that dynamically hedge against these climate news shocks. Similarly, Pástor et al. (2022) find that the returns on their green factor track unexpected increases in a climate-concern index. These findings suggest that climate news serves as a distinct state variable for asset pricing, and that

hedging climate-related shocks is important for understanding green premia. Relatedly, Choi et al. (2020) provide evidence that attention to global warming varies over time and is associated with return differences across carbon-intensive firms, supporting the view that climate attention is a relevant state variable for asset prices.

Capital flows and market microstructure can amplify or distort ESG signals. A key driver of time variation in ESG demand is the salience of third-party sustainability rankings. Hartzmark and Sussman (2019) exploit the introduction of Morningstar sustainability ratings (March 2016) and show that published sustainability rankings caused large reallocations of mutual fund flows across rated funds, consistent with discrete demand shocks tied to ESG ratings. For example, Beck (2021) shows that the performance of ESG investments has been largely driven by price pressure from investor fund flows. In his analysis, moving one dollar from the market portfolio into the representative ESG fund increases the aggregate value of high-ESG stocks by roughly 2 to 2.50. Without these sustained inflows, the ESG sector would underperform, and the positive alpha of a "green preference" portfolio would turn negative. Consistent with this flow-driven effect, index rebalancing events produce similar patterns. Rudkin et al. (2023) examine additions and deletions to the Dow Jones Sustainability Index (DJSI). They find significant positive pre-announcement abnormal returns when firms are added to the index, and those gains persist even after the index inclusion. Firms deleted from the DJSI also experience positive abnormal returns (as investors rebalance out of them), albeit with weaker significance. The index effect is especially pronounced for companies outside the S&P 500, suggesting that inclusion conveys more new information for lesser-known firms. However, the authors caution that conventional event-study methods may overstate these abnormal returns, and that the observed index inclusion effects likely reflect temporary demand shocks and microstructure frictions rather than changes in fundamental value.

Market liquidity is another channel that mediates observed green premia. Luo (2022), examining U.K. stocks, finds that firms in the lowest ESG-score quintile outper-

form those in the highest quintile by about 0.5% per month, but this return differential is significant only among less liquid stocks. He also shows that high-ESG firms tend to have tighter bid–ask spreads and deeper order books, and that their apparent alpha vanishes after controlling for liquidity. Moreover, many ESG sub-components (e.g., resource usage, emissions, workforce metrics) exhibit significant return relations, and companies with higher liquidity generally have higher ESG scores. In a Chinese market context, Chen et al. (2023) report that strong ESG performance improves a stock's liquidity by reducing agency costs, attracting greater foreign ownership, and enhancing firm reputation. This liquidity benefit is especially pronounced for state-owned enterprises and for firms in heavily polluting industries. Together, these microstructure-related findings suggest that superior ESG characteristics may lower expected returns via improved liquidity and reduced transaction costs, rather than via a fundamental risk factor commanding a premium.

Market crises provide natural experiments for examining ESG as a source of resilience. Albuquerque et al. (2020) analyze U.S. stock performance during the COVID-19 crisis (first quarter of 2020) and find that companies with higher environmental and social ratings had higher stock returns, lower return volatility, and higher operating profit margins during the crash. They attribute this resilience to strong customer and investor loyalty associated with ESG-oriented firms. However, evidence from other markets is mixed. Abedifar et al. (2022) report that while high-ESG firms in the United States experienced lower volatility during the COVID crash, comparable high-ESG firms in Japan and the U.K. actually underperformed both during and after the crash period. Meanwhile, Karolyi et al. (2025) highlight a representation bias in cross-country comparisons of green returns. They show that standard ESG datasets under-represent non-U.S. markets, which can distort inferences. After correcting for this bias, green-minus-brown return differentials of around 38 basis points per month are observed in developed markets outside the U.S.—comparable to the 41 basis points per month found within the U.S. These studies suggest that ESG premia can vary significantly by

region and sample period, and that cross-sectional heterogeneity must be taken into account when interpreting any estimated greenium.

Some recent work also attempts to reconcile evidence across different asset classes. Eskildsen et al. (2024) emphasize that if ESG preferences and risk factors are consistently integrated, the price of greenium should, in theory, be coherent across equities, corporate bonds, and sovereign bonds. They provide cross-asset evidence that the negative price of green beta indeed persists across these asset classes, and that the green discount is larger in countries with stronger sustainability preferences. Consistent with these findings, our own implied cost of capital regressions also indicate a small negative price of green beta, and in our ICC-based tests we find economically modest but statistically detectable state dependence. Our ICC-based evidence complements this literature by showing that climate-attention shocks change the *composition* of sustainability pricing. The exposure price becomes less negative after concern spikes, while the level component becomes more negative, consistent with attention shifting the balance between hedging-style exposure valuation and characteristic-based demand.

Taken together, prior work points to two distinct reasons green assets can command lower expected returns: investors may value ESG characteristics directly (a taste channel) and/or value exposure that hedges sustainability-state fluctuations (an exposure channel). At the same time, inference is complicated by rating disagreement and by flow- and liquidity-driven price pressure that can contaminate realised-return evidence (Berg et al., 2022; Christensen et al., 2022; Beck, 2021; Luo, 2022). These considerations motivate (i) using forward-looking expected-return proxies such as ICC, and (ii) measuring sustainability risk through an explicit ESG exposure, rather than relying on ESG levels alone.

3 Theory

3.1 Setting

The economy consists of $N+1$ assets with returns $\mathbf{r} = (r_1, \ldots, r_{N+1})^\top$, expected returns $\boldsymbol{\mu} = \mathbb{E}[\mathbf{r}]$, and covariance matrix

$$\Sigma = \mathbb{E}\Big[(\mathbf{r} - \boldsymbol{\mu})(\mathbf{r} - \boldsymbol{\mu})^\top\Big].$$

Asset $N+1$ is risk-free and yields return r^f. Each risky asset $n = 1, \ldots, N$ is associated with a green score s_n. We collect green scores in the vector

$$\mathbf{s} = (s_1, \ldots, s_N, 0)^\top,$$

where the risk-free asset has zero green exposure. Let $\bar{\mathbf{s}} = \mathbb{E}[\mathbf{s}]$ denote the vector of expected green scores and

$$\Omega = \mathbb{E}\Big[(\mathbf{s} - \bar{\mathbf{s}})(\mathbf{s} - \bar{\mathbf{s}})^\top\Big]$$

their covariance matrix.

There are I investors. Investor i has wealth share $w_i > 0$ and chooses portfolio weights $\mathbf{x}_i \in \mathbb{R}^{N+1}$. For any portfolio $\mathbf{x}$, define expected return and variance as

$$\mu_{r,\mathbf{x}} = \mathbf{x}^\top \boldsymbol{\mu}, \qquad \sigma^2_{r,\mathbf{x}} = \mathbf{x}^\top \Sigma \mathbf{x},$$

and expected green score and green-score variance as

$$\mu_{s,\mathbf{x}} = \mathbf{x}^\top \bar{\mathbf{s}}, \qquad \sigma^2_{s,\mathbf{x}} = \mathbf{x}^\top \Omega \mathbf{x}.$$

We extend the preference model of Pedersen et al. (2021) by allowing investors to care not only about the level of green exposure, but also about the risk associated with

the evolution of firms' green scores. Investor i maximizes

$$u_i(\mathbf{x}) = \mu_{r,\mathbf{x}} - \frac{1}{2\gamma_{r,i}}\,\sigma^2_{r,\mathbf{x}} + \eta_i\,\mu_{s,\mathbf{x}} - \frac{1}{2\gamma_{s,i}}\,\sigma^2_{s,\mathbf{x}}, \tag{3.1}$$

where $\gamma_{r,i} > 0$[1] denotes risk aversion with respect to financial returns, η_i captures the strength of green preferences, and $\gamma_{s,i} \neq 0$ governs preferences over green-score variability. We allow $\gamma_{s,i}$ to take negative values, reflecting that some investors may prefer portfolios with volatile green scores — for instance, because higher variability increases the likelihood of future improvements in firms' sustainability performance. [2]

3.2 Green Security Market Line

We analyse the asset pricing implications of the market model just described. At equilibrium each investor optimally solves the allocation problem of Equation (3.1) and choose portfolio $\mathbf{x}_i \in \mathbb{R}^{N+1}$. Moreover, market clears, i.e., $\tau_n = \sum_{i=1}^{I} \frac{w_i\, x_{i,n}}{1-x_{i,N+1}}$, where $\boldsymbol{\tau} \in \mathbb{R}^N$ is the market portfolio. For each investor i the first-order conditions imply, for assets $n = 1,\ldots,N+1$,

$$\mu_n + \eta_i\,\bar{s}_n = \frac{1}{\gamma_{r,i}}\,\mathrm{Cov}\left(r_n, \mathbf{x}_i^\top \mathbf{r}\right) + \frac{1}{\gamma_{s,i}}\,\mathrm{Cov}\left(s_n, \mathbf{x}_i^\top \mathbf{s}\right) + \theta_{\mathbf{x}_i}, \qquad \mathbf{1}^\top \mathbf{x}_i = 1. \tag{3.2}$$

For $n = N+1$ we obtain $\theta_{\mathbf{x}_i} = r^f$ for all $i = 1,\ldots,I$. The following result holds:

Theorem 1 (Green Security Market Line). *Suppose that* $\frac{\gamma_{r,i}}{\gamma_{s,i}} = \xi$ *for all* $i = 1,\ldots,I$. *It*

[1] We use a risk-tolerance parameterization; $1/\gamma_{r,i}$ corresponds to the standard mean–variance risk-aversion coefficient (as in Pedersen et al. (2021)).

[2] To ensure the portfolio problem is well-defined even when $\gamma_{s,i} < 0$, we assume that the quadratic term in (3.1) is concave (and hence the objective is bounded above) over the feasible set. Let Σ_R and Ω_R denote the $N \times N$ submatrices of Σ and Ω corresponding to the risky assets (excluding the risk-free asset). A sufficient condition is

$$A_i \;\equiv\; \frac{1}{\gamma_{r,i}}\,\Sigma_R \;+\; \frac{1}{\gamma_{s,i}}\,\Omega_R \;\succ\; 0,$$

which implies strict concavity in risky holdings and therefore a unique optimal portfolio under the budget constraint $\mathbf{1}^\top x_i = 1$. Under the aggregation restriction $\gamma_{r,i}/\gamma_{s,i} = \xi$, this condition is equivalent to $\Sigma_R + \xi\,\Omega_R \succ 0$.

follows that:

$$\mu_n - r^f = \left(\frac{1}{\bar{\gamma}_r}\,\sigma^2_{r,\boldsymbol{\tau}}\right)\beta_{n,\boldsymbol{\tau}} + \left(\frac{\xi}{\bar{\gamma}_r}\,\sigma^2_{s,\boldsymbol{\tau}}\right)\beta^g_{n,\boldsymbol{\tau}} - \frac{\bar{\zeta}_r}{\bar{\gamma}_r}\,\bar{s}_n, \tag{3.3}$$

where $\bar{\gamma}_r = \sum_{i=1}^{I} \frac{w_i \gamma_{r,i}}{1-x_{i,N+1}}$, $\bar{\zeta}_r = \sum_{i=1}^{I} \frac{w_i \eta_i \gamma_{r,i}}{1-x_{i,N+1}}$, *and*

$$\beta_{n,\boldsymbol{\tau}} = \frac{\mathrm{Cov}\,(r_n, \boldsymbol{\tau}^\top \mathbf{r})}{\sigma^2_{r,\boldsymbol{\tau}}}, \qquad \beta^g_{n,\boldsymbol{\tau}} = \frac{\mathrm{Cov}\,(s_n, \boldsymbol{\tau}^\top \mathbf{s})}{\sigma^2_{s,\boldsymbol{\tau}}}.$$

Theorem 1 requires that the ratio of financial to green risk tolerance,

$$\frac{\gamma_{r,i}}{\gamma_{s,i}},$$

is common across investors, while allowing full heterogeneity in green tastes η_i. This restriction has a clear economic interpretation: it assumes that investors differ in how much they *care* about greenness (captured by η_i), but not in how they *trade off green risk relative to financial risk*.

Indeed, in the first-order conditions (3.2), financial risk enters scaled by $1/\gamma_{r,i}$, while green-score risk enters scaled by $1/\gamma_{s,i}$. The ratio $\gamma_{r,i}/\gamma_{s,i}$ therefore governs the relative price an investor assigns to one unit of green-score variance compared to one unit of return variance. Imposing a common ratio means that, conditional on taking risk, investors agree on the *unit exchange rate* between financial risk and green risk, even though they may disagree substantially on the desirability of green exposure itself (through η_i).

There are several reasons why this assumption is economically plausible. First, green-score risk is typically assessed and constrained using the same institutional risk-management infrastructure as financial risk (e.g., volatility targets, tracking-error limits, or risk budgets). When both types of risk are managed within a unified framework, it is natural that their relative tolerance is technologically or institutionally pinned down, leading to similar $\gamma_{r,i}/\gamma_{s,i}$ across investors, even if their ESG objectives differ.

Second, many ESG investors face *exogenous constraints* on green-risk exposure—arising from regulation, fiduciary standards, or mandate design—that scale proportionally with their overall risk capacity. In such settings, an investor with higher financial risk tolerance (lower $1/\gamma_{r,i}$) is mechanically allowed to bear proportionally more green-score risk (lower $1/\gamma_{s,i}$), keeping the ratio $\gamma_{r,i}/\gamma_{s,i}$ approximately constant across the investor population.

Third, the restriction plays the role of a standard *aggregation condition*. With heterogeneous risk preferences, market-clearing prices generally depend on the full cross-sectional distribution of portfolios. Assuming a common $\gamma_{r,i}/\gamma_{s,i}$ aligns investors' marginal trade-offs between the two sources of risk, ensuring that individual demands aggregate into a tractable equilibrium characterized by a single green-risk price. Importantly, this does not restrict heterogeneity in tastes η_i, which continue to shape equilibrium through the taste-based term in the Green Security Market Line.

Finally, the assumption can be interpreted as a first-order approximation for the marginal price-setting investors. Even if $\gamma_{r,i}/\gamma_{s,i}$ varies across agents, as long as it is approximately constant among the wealth-weighted group of investors who determine prices, the equilibrium pricing implications of Theorem 1 remain valid. In this sense, the restriction is analogous to standard homogeneity assumptions routinely used to derive CAPM-style pricing relations.

4 Data

Our sample consists of U.S. common stocks from January 2007 to December 2023 (CRSP share codes 10 and 11). Following Jensen et al. (2023), we exclude nano- and micro-cap stocks (based on NYSE size breakpoints) to limit microcap- and liquidity-driven effects. ESG scores, implied costs of capital, returns, and firm characteristics are merged at the firm–month level using CRSP–Compustat identifiers. Unless stated otherwise, continuous variables are winsorised at the 1st/99th percentiles within month.

ESG measures

Our primary ESG metric is the MSCI ESG rating score obtained at a monthly frequency. MSCI ESG Ratings evaluate a company's sustainability performance on a 0–10 continuous scale. The MSCI score we use is a weighted-average ESG score that reflects the firm's performance across key ESG issues. For example, MSCI's Environmental Pillar Score (E-score) is measured on a 0–10 scale as a weighted average across 13 environmental issues (e.g., carbon emissions, water stress, etc.). We utilise the overall ESG score as our main measure of "greenness" and also consider the Environmental Pillar Score separately in robustness checks. All MSCI ESG data are industry-unadjusted scores (absolute performance). This ensures that a firm's score reflects its absolute sustainability performance rather than its rank within an industry, aligning with recent climate-finance research that constructs "brown" vs. "green" measures on an absolute basis. We winsorise the ESG scores to the 1st and 99th percentiles to limit the impact of any extreme rating values. When constructing the market ESG benchmark and estimating green betas, we use the winsorized MSCI scores on their original 0–10 scale. In the cross-sectional pricing tests, we rescale the ESG level and green-beta measures using fixed (time-invariant) cross-sectional standard deviations so that reported coefficients can be read as basis-point changes per one-standard-deviation difference and are comparable over time (see Section 5). In our sample, the average MSCI ESG score is 4.7, indicating moderate sustainability performance, with a standard deviation of roughly one point (See Appendix A.1). The Environmental pillar scores tend to be slightly lower on average than the Social or Governance pillars. Consistent with prior findings that environmental scores have the lowest mean among the ESG components, this reflects the greater challenges that firms face on environmental metrics.

We also construct an emissions-based greenness proxy to capture firms' direct carbon efficiency. Specifically, we use S&P Global Trucost greenhouse-gas (GHG) emissions intensity measures for Scope 1 (direct emissions) and Scope 2 (purchased electricity/steam/heat). Trucost reports these variables as emissions per unit of revenue

(typically measured in tons of CO_2-equivalent per $1 million of revenue). We sum Scope 1 and Scope 2 intensities to obtain a firm-level Scope 1+2 intensity measure relative to the firm's total assets.

To improve robustness given the strong right-skewness of carbon-intensity data, we winsorise the Trucost-based intensity measure at the 1st and 99th percentiles. Our main Trucost greenness characteristic is

$$\text{Greenness}_{i,t}^{\text{Trucost}} = -\text{Intensity}_{i,t}^{S1+S2,\text{A}},$$

so that higher values correspond to *lower* emissions intensity (i.e., a greener firm).

Trucost intensity is observed at an annual reporting frequency. To align it with our monthly panel, we map each annual observation to month-end dates and carry the same value forward for the subsequent 12 months (until the next annual update). This lower-frequency updating implies that the Trucost-based proxy is mechanically less time-varying than monthly MSCI scores and noisier at the monthly horizon. We use this emissions-based proxy as a complementary measure of greenness to assess whether pricing patterns are driven by direct carbon efficiency rather than broader ESG assessments.

Financial data and control variables

We obtain stock return and firm-characteristic data from the Global Stock Returns and Characteristics dataset of Jensen et al. (2023). This dataset (often referred to as the JKP global factors data) provides monthly stock-level data on returns and a large set of firm attributes for 93 countries, including the U.S. From this source we extract the U.S. subset and use key firm-level control variables analogous to those in prior studies. In particular, we control for characteristics known to predict returns ("factor zoo" variables): market beta, size (market capitalisation), value (e.g., book-to-market equity), momentum (past 12-month return), and short-term and long-term reversal

factors. These characteristics are constructed following standard definitions in the literature and are updated monthly in the JKP dataset. We use the characteristic data as provided, apart from aligning the timing (e.g., using lagged characteristics to predict returns). By including this broad set of controls, we aim to isolate the effect of ESG from other known return predictors. For example, controlling for size and value is important because ESG scores could be correlated with firm size or value/growth tendencies (larger firms and growth firms might have different ESG profiles). Our control variable choices mirror those of recent ESG asset-pricing studies, which typically include market beta, size, book-to-market and momentum in return regressions. All firm-level controls are winsorised at extreme values and, when used in regression, are standardised to have convenient units so that coefficients are comparable.

For portfolio-level analyses and robustness checks, we also incorporate factor return data from Kenneth French's data library. Specifically, we download the Fama–French factor returns for Market (MKT–RF), Size (SMB), Value (HML) and the Momentum (UMD) factor on a monthly basis. The five-factor specification adds profitability (RMW) and investment (CMA), so the spanning tests directly address the concern that green-beta premia may proxy for well-known sources of cross-sectional returns. The use of standard factor returns from French's library ensures our results can be compared to the wider asset-pricing literature.

Sentiment and ICC measures

In our analysis, we augment the above data with two supplementary measures to capture broader market forces. The first is a climate-concern index to proxy for time-varying investor sentiment or attention to climate issues. We use the Media Climate Change Concerns (MCCC) index developed by Ardia et al. (2021). This index is constructed via textual analysis of major U.S. newspapers, quantifying the daily intensity of news related to climate change. The MCCC index is available from 2003 onward; we use the monthly average values of this index (higher values indicate greater

public concern about climate change in that month). Intuitively, spikes in the MCCC index reflect heightened awareness of climate issues, which may drive "green" investor flows or repricing of green vs. brown stocks. Including MCCC as an explanatory variable helps us assess whether green asset returns co-move with shifts in climate concern. For example, Engle et al. (2020) and others find that on days of unexpected increases in climate news, green stocks outperform brown stocks. In our sample, we observe considerable time variation in MCCC. For example, we note surges around events like the Paris Agreement (2015) and an overall upward trend in the late 2010s, which we account for when analysing return differentials.

Our main expected-return proxy is the implied cost of equity capital (ICC), constructed following Eskildsen et al. (2024). Each month t, we compute four model-specific ICC estimates based on Gebhardt et al. (2001), Claus and Thomas (2001), Easton (2004), and Ohlson and Jüttner-Nauroth (2005). We winsorise each model-specific estimate at the 1st and 99th percentiles within month t and define $\text{ICC}_{i,t}$ as the equal-weighted average of the winsorised estimates. The ICC is an *annualized* discount rate observed at a monthly frequency.

Dataset construction

We merge MSCI ESG scores, ICC data and financial characteristics using the CRSP–Compustat link table. Our sample includes firm-month observations where all key variables (ESG score, ICC, returns and control characteristics) are available. ESG scores and firm characteristics are aligned to month-end. The implied cost of capital(ICC) is dated t and merged contemporaneously to the firm-month panel. Green beta estimation requires a 36-month rolling window; combined with months in which ICC coverage is missing, the final estimation sample contains 152 monthly cross-sections. We align all realized-return tests to the same set of month-end dates for comparability.

5 Methodology

The empirical design estimates how exposure to market-wide ESG shifts maps into required returns. Let i index firms and t months. Denote by $P_{i,t}$ the price, by $R_{i,t+1}$ the gross return from t to $t+1$, and by $R_{m,t+1}$ the corresponding market return. We write excess returns as $r_{i,t+1} = R_{i,t+1} - R_{f,t+1}$ and $r_{m,t+1} = R_{m,t+1} - R_{f,t+1}$.[3]

Raw ESG scores $es_{i,t}$ are aligned to month-end and winsorised at the 1st/99th percentiles within month t. We use these *winsorized* ESG levels to construct the market ESG benchmark and estimate firm-level green betas. In the cross-sectional pricing tests, we express greenness effects per one-standard-deviation difference using a *fixed* (time-invariant) scale so that coefficients (and the implied monthly price series) are comparable over time and not mechanically affected by changes in cross-sectional dispersion. Specifically, we define

$$s_{i,t} = \frac{es_{i,t} - \overline{es}_t}{\bar{\sigma}_{es}}, \qquad \bar{\sigma}_{es} \equiv \frac{1}{T}\sum_{t=1}^{T} \sigma_t(es), \tag{5.1}$$

where $\overline{es}_t$ and $\sigma_t(es)$ denote the cross-sectional mean and standard deviation of $es_{i,t}$ in month t. Unlike within-month z-scoring, this fixed-scale convention ensures that time variation in estimated monthly prices is not mechanically driven by changes in cross-sectional dispersion. Equivalently, one can estimate the monthly cross-sections using demeaned winsorized ESG levels and then rescale the resulting slope by $\bar{\sigma}_{es}$ for reporting; we use the scaled-regressor notation for transparency.[4]

Market ESG score. The monthly market ESG benchmark is the value-weighted cross-sectional aggregate of ESG scores,

$$s_{m,t} = \sum_{j \in U_t} w_{j,t}\, es_{j,t}, \qquad w_{j,t} = \frac{ME_{j,t}}{\sum_{k \in U_t} ME_{k,t}},$$

[3] When the dependent variable is one-month-ahead realized excess returns $r_{i,t+1}$, all right-hand-side variables are dated t (portfolio formation) and returns are realized over $t \to t+1$.

[4] Results are robust to alternative fixed-scale conventions, including using the full-sample standard deviation and robust scaling based on the full-sample median and MAD.

where U_t is the investable universe in month t and $ME_{j,t}$ is market equity.[5]

For each firm i and month t, we estimate green beta $\beta^g_{i,t}$ from the ESG–ESG sensitivity, using a rolling window of $L = 36$ months with a minimum of 12 observations:

$$\text{(Green beta)} \qquad es_{i,\tau} = \alpha^g_{i,t} + \beta^g_{i,t} s_{m,\tau} + u_{i,\tau}, \qquad \tau \in \{t - L + 1, \ldots, t\}. \tag{5.2}$$

Estimates of $\widehat{\beta}^g_{i,t}$ are stored at month t. We require at least 12 non-missing pairs $(es_{i,\tau}, s_{m,\tau})$ for the rolling regression; otherwise $\widehat{\beta}^g_{i,t}$ is set to missing. To reduce the impact of small-sample noise, we winsorise $\widehat{\beta}^g$ each month at the 1st/99th percentiles.[6]

Equation (5.2) defines the green beta $\beta^g_{i,t}$ as the sensitivity of a firm's ESG score to movements in the market ESG benchmark $s_{m,t}$. In the model, $s_{m,t}$ plays the role of the greenness state variable that enters the Green Security Market Line in Theorem 1. Since $\beta^g_{i,t}$ is constructed from ESG scores rather than returns, it should be interpreted as an exposure measure and not as a loading on a traded return factor.

Green beta is not a return beta. It measures how a firm's *perceived* sustainability, captured by its ESG score, co-moves with aggregate shifts in market greenness. A high-β^g firm is one whose ESG profile tends to strengthen when aggregate ESG rises (and, more generally, in periods when sustainability becomes more salient), whereas a low-β^g firm moves little with the aggregate series. Within the Green Security Market Line framework, β^g therefore captures *exposure* to a (non-traded) sustainability state variable rather than a loading on a traded return factor. We therefore use ICC, a forward-looking proxy for expected returns, as the primary pricing target and view realized-return specifications as complementary diagnostics. A negative price of β^g means investors accept lower required returns for holding assets that deliver favorable exposure to that state (e.g. assets that tend to look greener precisely when sustainability is most valued in markets). This distinction is central: ESG *levels* summarize a firm's

[5] An equal-weighted alternative is used in robustness. We also consider first differences $\Delta es_{i,t} = es_{i,t} - es_{i,t-1}$ and $\Delta s_{m,t}$ in robustness; the baseline uses levels.

[6] In thin windows we also consider ridge-stabilized estimation; results are unchanged.

standing at a point in time, while green beta summarizes its sensitivity to systematic sustainability movements.

We estimate prices of market risk, green-beta exposure, and ESG level in the cross-section via Fama–MacBeth two-step regressions (Fama and MacBeth, 1973). Let $Y_{i,t}$ denote the left-hand-side pricing target. We use either (i) a forward-looking implied cost of capital (ICC) measured at t, denoted $\mathrm{ICC}_{i,t}$, or (ii) one-month-ahead realized excess returns $r_{i,t+1}$. For each month t, we run the cross-sectional regression

$$Y_{i,t} = a_t + \lambda_{g,t}\,\widetilde{\beta}^{g}_{i,t} + \lambda_{m,t}\,\widehat{\beta}^{m}_{i,t} + \pi_{s,t}\,s_{i,t} + \theta_t' X_{i,t} + \varepsilon_{i,t}. \tag{5.3}$$

where $X_{i,t}$ includes the controls used in our main tables: size, value, and the past-return controls. Cross-sections are equal-weighted (value-weighted robustness reported). We rescale $\widehat{\beta}^{g}_{i,t}$ using the same fixed-scale convention,

$$\widetilde{\beta}^{g}_{i,t} = \frac{\widehat{\beta}^{g}_{i,t} - \overline{\beta}^{g}_{t}}{\bar{\sigma}_{\beta^g}}, \qquad \bar{\sigma}_{\beta^g} \equiv \frac{1}{T}\sum_{t=1}^{T} \sigma_t(\widehat{\beta}^{g}),$$

so that coefficients can be interpreted as changes in annual ICC (or returns) per one standard deviation of green-beta exposure on a constant scale. The estimated monthly slopes $\{\widehat{\lambda}_{g,t}, \widehat{\pi}_{s,t}\}$ can be summarized by their time-series averages, $\bar{\lambda}_g = (1/T)\sum_{t=1}^{T}\widehat{\lambda}_{g,t}$ and $\bar{\pi}_s = (1/T)\sum_{t=1}^{T}\widehat{\pi}_{s,t}$. Inference uses Newey–West HAC standard errors with $q = 6$ lags. For realized-return specifications, we also report Shanken-adjusted t-statistics to account for generated regressors (Shanken, 1992).

Computation of implied cost of capital. In our analysis, we use a forward-looking proxy for expected returns: the *implied cost of capital* (ICC). Using ICC mitigates transitory noise in realized returns and reduces contamination from short-lived taste shocks or flow-driven price pressure. Conceptually, ICC is the internal rate $k_{i,t}$ that

equates observed price to the present value of expected cash flows:

$$P_{i,t} = \sum_{h=1}^{H} \frac{\mathbb{E}_t[CF_{i,t+h}]}{(1+k_{i,t})^h} + \frac{\mathbb{E}_t[TV_{i,t+H}]}{(1+k_{i,t})^H}, \tag{5.4}$$

where cash flows CF are dividends or model-implied payouts, and TV is a terminal value. We implement a standard set of ICC estimators widely used in the literature and also referenced in Eskildsen et al. (2024): (i) the residual-income model of Gebhardt et al. (2001), which forecasts ROE mean reversion and solves for k in the residual-income identity; (ii) the dividend/earnings-growth model of Claus and Thomas (2001); (iii) the PEG estimator of Easton (2004), which infers k from near-term analyst EPS growth relative to price; and (iv) the abnormal-earnings-growth (OJ) model of Ohlson and Jüttner-Nauroth (2005). For each firm-month t, we:

1. collect prices, book value of equity, and analyst EPS/dividend forecasts (where available);

2. form cash-flow or earnings paths over a finite horizon H (typically 3–5 years), with payout linked to earnings via a dynamic payout ratio;

3. set a terminal value using a long-run growth anchor (inflation-plus-real growth) and mean-reverting profitability assumptions (for residual-income);

4. solve numerically for $k_{i,t}$ that satisfies (5.4) for each estimator; and

5. winsorise each model-specific ICC estimate cross-sectionally within month t at the 1st and 99th percentiles and define the firm-month ICC as the equal-weighted average across the available model-specific estimates.

Our ICC series is constructed and observed at the monthly frequency, using prices and the most recent available cash-flow expectations as of month t. Following standard practice in the ICC literature, we aggregate multiple ICC estimators to reduce model-specific noise and obtain a more robust forward-looking proxy for expected returns.

To ease interpretation, we report greenness slopes in annual basis points (multiplying ICC-unit coefficients by 10,000) per one-standard-deviation difference using the fixed scaling defined above.

All right-hand-side variables are known at t. We drop firm-months with missing $\hat{\beta}^{g}_{i,t}$ or $\hat{\beta}^{m}_{i,t}$, or with extreme values in ICC specifications. We adjust for share code changes, require CRSP/Compustat linkages to be valid, and remove microcaps below the 20th percentile of ME in robustness. These construction steps yield a monthly panel of implied costs of capital, estimated green betas, ESG score levels, and standard firm characteristics, aligned so that all right-hand-side information is available at the month in which discount rates are measured.

6 Results

This section evaluates whether systematic exposure to market-level greenness carries a distinct price from firms' ESG levels, and how both components vary over time. Our primary pricing tests use the implied cost of capital (ICC), which is a forward-looking proxy for expected returns and therefore provides sharper inference than short-horizon realized-return spreads.

We proceed in five steps. Section 6.1 establishes stock-level pricing of green beta in ICC cross-sections while controlling for ESG levels and standard characteristics. Section 6.2 studies whether the prices of the exposure and level channels vary with shocks to climate concern (MCCC). Section 6.3 documents longer-run time variation in both channels. Section 6.4 asks whether the green-beta relation reflects industry tilts or within-industry dispersion, using FF49 industry portfolios and within–between decompositions. Section 6.5 summarizes portfolio implications.

6.1 Stock-level pricing of green beta and the ESG level

We begin with cross-sectional pricing tests based on ICC. Each month t, we estimate month-by-month cross-sectional regressions of firm-level ICC on market beta, green beta, the ESG score level, and standard characteristics and summarize the time-series of slope estimates following Fama–MacBeth regressions. The benchmark specification is

$$\mathrm{ICC}_{i,t} = a_t + \lambda_{m,t}\beta^m_{i,t} + \lambda_{g,t}\widetilde{\beta}^g_{i,t} + \pi_{s,t}s_{i,t} + X'_{i,t}\theta_t + \varepsilon_{i,t}, \tag{6.1}$$

where $\widetilde{\beta}^g_{i,t}$ is the firm-level green beta and $s_{i,t}$ is the ESG score level. The greenness regressors are centered each month and rescaled by fixed standard deviations (Section 5), so $\lambda_{g,t}$ and $\pi_{s,t}$ can be read as changes in annual ICC associated with a one-standard-deviation difference in exposure/level on a constant scale.

Green beta is priced after controlling for ESG level. Table 1 summarizes the monthly Fama–MacBeth cross-sections. The main pattern is that firms whose ESG profiles comove more strongly with aggregate market greenness exhibit systematically lower implied discount rates. In the preferred specification with the full set of controls (Spec. D), the green-beta slope is negative and statistically significant, with an estimate of -0.0012 (Newey–West $t = -2.39$). Since $\widetilde{\beta}^g_{i,t}$ is scaled by a fixed cross-sectional standard deviation, this corresponds to about -12 annual basis points in ICC per one-standard-deviation higher green beta. The estimate is stable across the nested specifications and remains negative after conditioning on the ESG score level and the standard characteristic set. The persistence of λ_g alongside π_s is the central empirical message of the paper: green beta captures a priced *exposure* dimension that is not subsumed by static ESG classifications. Appendix Table A.13 shows that this result is not driven by short-horizon ESG score revisions: the baseline (level) green beta remains negatively priced when we jointly include transitory proxies based on ESG changes or AR(1)-filtered innovations.

The ESG score level is also priced, and more precisely so. In the full specification, a one-standard-deviation higher ESG score is associated with an estimate of −0.0023 (Newey–West $t = -8.88$), corresponding to about −23 basis points lower annual ICC. This pattern is consistent with a valuation wedge tied to ESG characteristics, as emphasized in preference-based frameworks in the literature, and it lines up with broader ICC-based evidence that the greenium is economically modest but not negligible (Pástor et al., 2022; Pedersen et al., 2021; Eskildsen et al., 2024). The joint significance of λ_g and π_s in our results suggests that the empirical greenium is not a single object. Expected discount rates reflect both pricing of *levels* (ESG characteristics) and an incremental price for *systematic sustainability exposure* measured by green beta.

Economically, the estimates imply a small but systematic wedge in firms' costs of equity. In Spec. D, a one-standard-deviation higher green-beta exposure is associated with about 12 annual basis points lower ICC, holding the ESG level and standard characteristics fixed, while a one-standard-deviation higher ESG level is associated with an additional 23 basis point reduction. A firm that is one standard deviation higher along both dimensions therefore has a predicted ICC that is roughly $12 + 23 \approx 35$ bps lower per year. While modest relative to the level of equity discount rates, wedges of this size are precisely the effects that are most visible in forward-looking expected-return proxies such as ICC (Eskildsen et al., 2024). Interpreted through the GSML, the ESG-level slope captures a valuation wedge tied to static sustainability characteristics, while the green-beta slope captures pricing of co-movement with the aggregate greenness state.

Realized-return tests are less precise at monthly horizons. When we repeat the cross-sectional tests using one-month-ahead realized excess returns, the estimated greenness slopes are substantially noisier and statistically indistinguishable from zero. Appendix Table A.4 reports the corresponding two-pass estimates.

If sustainable investing operated purely through *taste* for high-ESG firms, then once we control for the ESG level $s_{i,t}$ (and standard characteristics) the exposure coefficient

$\lambda_{g,t}$ should be statistically indistinguishable from zero, while $\pi_{s,t} < 0$ should remain. Conversely, under a pure *risk/hedging* view in which investors price only systematic sustainability exposure, $\lambda_{g,t}$ should remain negative while the level slope $\pi_{s,t}$ should be small. The evidence in Table 1 rejects both corner cases: both $\beta^g_{i,t}$ and $s_{i,t}$ enter with economically meaningful and statistically significant slopes in the full specification. Sections 6.2–6.3 further show that the two prices move differently with climate salience and over time, reinforcing that exposure and level capture distinct economic channels.

6.2 Climate concern and state-dependent pricing of the two channels

Next, we examine whether the prices of green-beta exposure and ESG tastes vary with the climate state. We extract monthly price series $\{\lambda_{g,t}, \pi_{s,t}\}$ from the ICC cross-sections and relate them to innovations in climate concern measured by the Media Climate Change Concern (MCCC) index. Specifically, we estimate

$$\lambda_{g,t} = \bar{\lambda}_g + \delta_\lambda \, Shock_{t-1} + u_t, \tag{6.2}$$

$$\pi_{s,t} = \bar{\pi}_s + \delta_\pi \, Shock_{t-1} + v_t, \tag{6.3}$$

where $Shock_{t-1}$ is the standardised MCCC shock (lagged one month). Inference uses Newey–West corrections to accommodate serial dependence in the estimated monthly prices.

Table 2 relates the month-by-month sustainability prices extracted from the ICC cross-sections to the lagged innovation in climate concern. The estimates point to a reallocation across channels. After a positive climate-concern shock, the price of green-beta exposure rises toward zero, implying a narrower exposure-related discount. At the same time, the ESG-level price becomes more negative, indicating a wider level-based wedge. Because prices are reported in annual basis points per one-standard-deviation difference (Section 5), the Shock_{t-1} coefficients can be read as the change in each sustainability price associated with a $+1\sigma$ climate-attention shock.

Quantitatively, a one-standard-deviation increase in the lagged MCCC shock increases the price of green beta by 7.39 annual basis points (Newey–West $t = 2.87$). Since the average price is about −12 bps per one-standard-deviation exposure, the shock moves the exposure price materially toward zero. The score-level channel shifts in the opposite direction: the same shock decreases the ESG-level price by 2.67 bps (Newey–West $t = -1.80$), implying a modestly stronger (more negative) level discount in high-concern months.

These opposite-signed responses reject single-channel stories. Under a pure taste view, climate-concern shocks should mainly steepen the level wedge with little systematic movement in the exposure price. Under a pure hedging/risk view, the exposure price should move but the level wedge should not. Instead, the data show a state-dependent *reweighting* of the two components: climate attention shifts the balance between systematic greenness exposure and static ESG characteristics in discount-rate determination.

Appendix Tables A.9–A.12 verify that this conclusion is not sensitive to timing conventions or functional form. The exposure response is robust across specifications, whereas the score-level response is concentrated in nonlinear/high-concern regimes.

Under pure taste, climate-concern shocks should primarily steepen the level slope $\pi_{s,t}$ with little systematic effect on $\lambda_{g,t}$; under pure hedging/risk, the opposite. Table 2 shows that both channels respond, but in opposite directions: climate concern attenuates the green-beta discount (moves $\lambda_{g,t}$ toward zero) while strengthening the ESG-level discount (more negative $\pi_{s,t}$). This pattern is inconsistent with single-channel explanations and supports the interpretation that climate salience shifts the relative pricing of sustainability *levels* and *exposures*.

6.3 Time dependence: secular evolution in green-beta pricing and tastes

Beyond short-run climate-salience shocks, the pricing of sustainability may evolve as sustainable investing grows, measurement improves, and climate risks enter investor expectations more prominently. To quantify these longer-run dynamics, we test for linear trends in the monthly pricing objects and in the implied green-beta premium between high- and low-exposure portfolios.

Table 3 reports time-trend regressions for the green-beta price component, the score/taste component, the implied $Q5-Q1$ premium, and the cross-sectional dispersion in green betas. These results highlight two patterns. First, the taste-related price becomes progressively more negative, consistent with a strengthening discount-rate wedge tied to ESG levels. Second, the green-beta price displays a positive trend, indicating that the exposure channel evolves differently from the level channel. The decomposition therefore sharpens interpretation that secular changes in the "greenium" reflect the net outcome of two channels that can move in different directions.

The positive time trend in $\lambda_{g,t}$ indicates that the equilibrium valuation of green-beta exposure is not a structural constant. Appendix Table A.8 makes this shift transparent: the mean exposure price rises by about 31 annual bps (from −27 bps pre-2016 to about +4 bps post-2016), while the level/taste price $\pi_{s,t}$ becomes substantially more negative over the same split (from about −10 to about −31 bps). In the GSML framework, $\lambda_{g,t}$ is pinned down by the wealth-weighted set of marginal, price-setting investors and can therefore rotate as the salience of ESG metrics in delegated portfolio choice, mandates, and capital allocation changes over time. This interpretation is especially plausible around 2016, when third-party sustainability ratings became a prominent input into mutual-fund flows and investor attention (Hartzmark and Sussman, 2019), and when institutional investors report treating climate risks as financially material and addressing them through risk management and engagement (Krueger et al., 2020).

Consistent with a broader post-2016 regime shift in climate-related pricing, Gormsen et al. (2024) document a pronounced divergence in perceived green versus brown costs of capital since 2016. Taken together, these developments motivate interpreting the post-2016 compression of $\lambda_{g,t}$ as part of a broader "rotation" in sustainability pricing, in which the characteristic-based channel strengthens while the incremental valuation of co–movement with aggregate market greenness becomes weaker.

Why dispersion matters for portfolio premia. The cross-sectional dispersion in green betas increases over time, which mechanically amplifies the portfolio-relevant impact of the exposure price. Even if the per-unit price of green beta is modest, a widening spread in exposures implies that the implied premium between high and low green-beta portfolios can grow. Appendix Table A.8 provides a complementary split-sample illustration around 2016 and shows economically meaningful shifts in both $\lambda_{g,t}$ and the implied $Q5-Q1$ premium.

Because these estimates are obtained from firm-level cross-sections, they may reflect both systematic differences across industries and heterogeneity in green betas across firms within the same industry. The next subsection separates these channels using FF49 industry portfolios and within–between decompositions.

To verify that our findings are not specific to the aggregate ESG metric, we repeat the stock-level ICC tests using MSCI's Environmental Pillar Score both as the greenness characteristic and in the construction of the green factor used to estimate firms' green betas. The results confirm the central distinction between exposure and level. The environmental green beta remains negatively priced in the ICC cross-section and retains statistical significance across specifications, including when controlling for the environmental score level and the full set of firm characteristics. In contrast, the environmental score level exhibits limited incremental explanatory power once the controls are included. Appendix A.7 reports the corresponding ICC regressions and the associated time-series evidence on state-dependent pricing with climate-concern

shocks.

If the "greenium" reflected only tastes for ESG levels, secular dynamics should appear mainly in $\pi_{s,t}$ and in score-based premia, with little independent movement in $\lambda_{g,t}$. Conversely, if it reflected only priced sustainability risk/hedging, trends should primarily show up in $\lambda_{g,t}$ and exposure dispersion. The opposing trends in Table 3, a stronger level discount alongside a gradual attenuation in the green-beta price, are inconsistent with either single-channel explanation and support the view that exposure and level are distinct priced objects.

6.4 Industry portfolios and the within–industry source of green-beta pricing

The stock-level evidence in Table 1 raises a natural question: does the pricing of β^g primarily reflect (i) *industry tilts* (persistent differences between greener and browner sectors) or (ii) *within-industry* dispersion in green betas across firms? We address this question in two complementary ways.

Let k index FF49 industries. For each month t, we form a value-weighted industry portfolio and compute its implied cost of capital as the value-weighted mean of firm-level ICCs within the industry. The resulting industry panel contains up to 44 industries per month in the ICC sample.

Leave-one-industry-out (LOO) industry green betas. At the industry level, estimating green betas relative to the aggregate market ESG benchmark raises a mechanical self-inclusion concern: the market ESG series contains industry k itself. We therefore estimate *leave-one-industry-out* industry green betas by regressing the value-weighted industry ESG score $s_{k,\tau}$ on a market ESG benchmark $s_{m,\tau}^{(-k)}$ computed excluding industry k, using a 36-month rolling window (minimum 12 observations), mirroring (5.2).

Table 4 reports Fama–MacBeth regressions run across industries each month. Specs (A)–(C) show that the industry-level price of green-beta exposure is negative and statistically

significant in parsimonious specifications. The coefficient on $\beta^g_{k,t}$ is -0.0033 ($t = -3.42$) in Spec. A and remains negative after controlling for market beta (Spec. B: -0.0026, $t = -2.97$) and in Spec. C (-0.0018, $t = -2.55$). After scaling by the cross-sectional dispersion of industry green betas , the estimates imply that a one-standard-deviation higher industry green beta is associated with a lower ICC of roughly 32 bps (Spec. A), 25 bps (Spec. B), and 17 bps (Spec. C).

In the fully controlled specification (Spec. D), the green-beta coefficient becomes economically small and statistically indistinguishable from zero: the point estimate implies about 7 bps per 1σ higher $\beta^g_{k,t}$. This difference does not contradict the firm-level evidence in Table 1. Industry portfolios are a coarse aggregation that mechanically removes within-industry dispersion in β^g; accordingly, once standard industry characteristics are included, little independent variation in industry-average β^g remains to be priced. Consistent with this interpretation, adding size alone in Columns (5)–(6) does not eliminate the negative green-beta coefficient; rather, the attenuation in Spec. D is driven by the combined inclusion of the value and past-return controls. Finally, the ESG *level* continues to be priced in the industry cross-section: with $\sigma(s_{k,t}) = 0.661$, the ESG coefficient in Specs. C–D corresponds to about 31 bps (Spec. C) and 27 bps (Spec. D) lower ICC per one-standard-deviation higher industry ESG level.

Within- vs. Between-industry decomposition at the firm level. To reconcile the strong firm-level pricing of β^g in Table 1 with the weaker incremental industry-level effect in saturated industry specifications, we decompose firm-month green beta into a

[6]To express coefficients in "per one cross-sectional standard deviation" units, we use fixed dispersion constants. For any industry-month variable $x_{k,t}$ (e.g., industry green beta or industry ESG), let $\sigma_t(x)$ denote the cross-sectional standard deviation across industries in month t (after winsorization and requiring at least 10 non-missing industries). We define $\bar{\sigma}_x \equiv \frac{1}{T}\sum_{t=1}^{T}\sigma_t(x)$ and report slopes multiplied by $\bar{\sigma}_x$. This is algebraically equivalent to estimating the cross-sectional regression using the standardized regressor $(x_{k,t} - \bar{x}_t)/\bar{\sigma}_x$.

within-industry component and a between-industry component:

$$\beta^g_{i,t} = \underbrace{\bar{\beta}^g_{k,t}}_{\text{between (industry mean)}} + \underbrace{\left(\beta^g_{i,t} - \bar{\beta}^g_{k,t}\right)}_{\text{within (deviation)}},$$

where k denotes the FF49 industry of firm i. Table 5 shows that the pricing relation is primarily a within-industry phenomenon. We show that the within-industry deviation is negative and statistically significant, while the between-industry component is small and statistically insignificant, both under equal-weighted and value-weighted industry means. Economically, a one-standard-deviation increase in the within-industry component implies roughly a 10–11 basis point reduction in ICC. Taken together, the evidence indicates that green-beta pricing in expected returns arises mainly from heterogeneity *within* industries rather than from a pure industry tilt, helping explain why industry-portfolio regressions lose incremental explanatory power once standard industry characteristics are included.

Within-industry neutral green-beta portfolios. To provide a nonparametric complement to the regression decomposition in Table 5, we form within-industry neutral green-beta portfolios that hold industry exposures fixed. Specifically, in each month t and within each FF49 industry k, we sort firms into quintiles based on the within-industry deviation $\beta^g_{i,t} - \bar{\beta}^g_{k,t}$ and compute quintile outcomes within each industry. We then aggregate the industry-level quintile outcomes either equally across industries or by industry market capitalization. By construction, the resulting quintile portfolios have identical industry weights each month; the $Q5-Q1$ spread therefore isolates within-industry variation in β^g rather than an industry tilt.

Table 6 reports time-series averages of the within-industry neutral portfolios for both realized one-month-ahead excess returns and implied costs of capital (ICC). The ICC spread remains negative after imposing industry neutrality. Across alternative weighting schemes, the $Q5-Q1$ ICC spread ranges from −0.0030 to −0.0039 (i.e., about

−30 to −39 annual bps). In the baseline EW/cap construction (Column (3)), the $Q5-Q1$ ICC spread equals −0.0030 (i.e., −30.0 bps; $t = -2.84$). Realized returns exhibit the same negative ordering but are noisier: in Column (3), the $Q5-Q1$ return spread is −0.0015 (i.e., −15.0 bps per month; $t = -1.41$).

Taken together with the firm-level cross-sections in Table 1, this portfolio evidence reinforces the results of Table 5. Green-beta pricing reflects heterogeneity *within* industries rather than reallocation across industries.

6.5 Portfolio implications: A factor-mimicking GB_t factor and spanning tests

Finally, we translate the stock-level pricing evidence into standard factor language by constructing a factor-mimicking *Green-Beta* return, GB_t. The factor is defined each month as the high–minus–low return spread between firms with high and low green-beta exposure. Because the preceding results show that green-beta pricing primarily reflects *within*-industry dispersion rather than an industry tilt, our baseline factor construction imposes industry neutrality. In each month t and within each FF49 industry, we sort firms into green-beta quintiles and compute the within-industry return spread between the highest and lowest quintiles. We then aggregate these industry-level spreads across industries ("spread aggregation"), which holds industry weights fixed by construction and isolates the within-industry component of β^g. [7]

Table 8 reports spanning regressions of the industry-neutral GB_t on standard Fama–French factors (FF3 and FF5) and momentum. The factor earns a negative average return of −24.8 bps per month ($t = -2.15$). Importantly, the intercept remains between −23.5 and −24.0 bps per month after controlling for FF3, FF5, and momentum (t-statistics around −2.3 to −2.4). Factor loadings are economically small and statistically

[7] Under the baseline weighting (VW within industries and equal-weighting across industries), the spread-aggregation procedure produces the same $Q5-Q1$ return series shown in Table 6 (column "VW / equal"); this series is the GB_t factor used in the spanning regressions. Portfolio formation requires at least 35 firms per industry-month, so months that fail this coverage requirement are omitted (145 months for realized returns; 151 months for ICC).

indistinguishable from zero, and the models explain little of the factor's variation ($R^2 < 0.02$). The spanning regressions show that the Green-Beta factor is not subsumed by conventional equity factors. It is not a re-labeled value, momentum, or generic market exposure.

The negative sign of GB_t is also consistent with the paper's central interpretation of green beta as a hedging/diversification attribute: firms with higher β^g command lower expected returns, so a long position in high-green-beta firms is associated with a lower realized return premium. For completeness, Appendix Tables A.6 and A.7 report analogous spanning regressions for alternative factor constructions (an unconditional sort and an alternative industry-neutral aggregation). While the magnitude varies across constructions, the sign of the Green-Beta premium is stable and the factor remains largely orthogonal to standard equity factors.

The evidence points to two priced sustainability components in expected returns. Firms with higher ESG levels have lower implied costs of capital, consistent with a preference-based discount. Separately, green beta remains negatively priced after controlling for ESG levels and standard characteristics, indicating that systematic sustainability exposure is priced in its own right. In the time series, climate-attention shocks tilt sustainability pricing toward the level component and away from the exposure component. This reinforces the central message that expected-return differences associated with greenness reflect two distinct objects whose relative importance varies with the attention state.

7 Conclusion

We propose a Green Security Market Line that separates sustainability pricing into a level component (ESG characteristics) and an exposure component (comovement with aggregate market greenness). Using U.S. equities and implied costs of capital as forward-looking discount-rate proxies, we find that both higher ESG levels and higher

green beta are associated with lower required returns. The joint significance of levels and exposures supports a two-channel interpretation in which investor preferences and hedging/diversification motives both contribute to the equilibrium pricing of sustainability.

The estimated slopes line up with the model's decomposition: the green-beta price is negative, consistent with investors valuing exposure to sustainability states, and the ESG-level slope is negative, consistent with a valuation wedge for high-ESG characteristics. Economically, the implied cost-of-capital differences are modest but persistent, on the order of tens of basis points per year between otherwise comparable greener and browner firms.

The evidence implies that greener firms can face lower required returns through both ESG-level discounts and priced exposure to aggregate ESG shifts, potentially incentivizing climate-friendly projects, as hypothesized in policy and stewardship discussions. At the same time, the relatively small magnitude implies that the greenium, while statistically significant, may not be large enough on its own to drastically alter capital allocation. This conclusion is also argued by Eskildsen et al. (2024) and others who argue the greenium is too small to replace policy tools like carbon taxes. Finally, climate attention shifts the composition of sustainability pricing. Following spikes in climate concern, the exposure-related price moves upward, reducing the absolute green-beta discount, while the level-related wedge becomes more negative. This nuance refines the narrative from earlier studies that climate news can generate episodic return patterns, but ex ante sustainability pricing reflects the interaction of a characteristic-based discount and a systematic exposure component whose relative weight varies across attention states.

A natural next step is to sharpen identification of the exposure channel. Because our green beta is estimated from the time-series comovement of provider scores, an important challenge is to separate economically meaningful sustainability shocks from slow-moving score dynamics and provider revisions. One approach is to estimate β^g

using innovations in market greenness (e.g., changes or residualized scores) and then benchmark the resulting exposures against discrete transition events and outcome-based anchors where feasible. A second extension is external validity across markets: applying the GSML level–exposure decomposition to corporate credit and sovereign debt would test whether the same two channels map into spreads and yields. Finally, the two-channel structure suggests a clientele test that shifts in ESG-mandated ownership, constraints, and flows should primarily move the level (taste) slope, while the exposure slope should be more closely tied to hedging demand (Hens and Trutwin, 2024).

References

Abedifar, Pejman, Kais Bouslah, Christopher Neumann, and Amine Tarazi (2022). "Resilience of Environmental and Social Stocks under Stress: Lessons from the COVID-19 Pandemic". In: *Financial Markets, Institutions & Instruments*. DOI: 10.1111/fmii.12166.

Albuquerque, Rui, Yrjö Koskinen, Shuai Yang, and Chendi Zhang (2020). "Resiliency of Environmental and Social Stocks: An Analysis of the Exogenous COVID-19 Market Crash". In: *The Review of Corporate Finance Studies* 9(3), pp. 593–621. DOI: 10.1093/rcfs/cfaa011.

Ardia, Nikos, Daniela Rapoport, and Peter Klaassen (2021). *Media Climate Change Concerns (MCCC) Index*.

Avramov, Doron, Si Cheng, Abraham Lioui, and Andrea Tarelli (2022). "Sustainable Investing with ESG Rating Uncertainty". In: *Journal of Financial Economics* 145(2), pp. 642–664. DOI: 10.1016/j.jfineco.2021.09.009.

Beck, Philippe van der (2021). "Flow-Driven ESG Returns". In: *Swiss Finance Institute Research Paper* (21-71).

Berg, Florian, Julian F. Kölbel, and Roberto Rigobon (2022). "Aggregate Confusion: The Divergence of ESG Ratings". In: *Review of Finance* 26(6), pp. 1315–1344. DOI: 10.1093/rof/rfac031.

Bolton, Patrick and Marcin Kacperczyk (2021). "Do Investors Care About Carbon Risk?" In: *Journal of Financial Economics* 142(2), pp. 517–549. DOI: 10.1016/j.jfineco.2021.05.008.

Chen, Ruipeng, Yirong Liu, and Wen Liu (2023). "Does ESG Performance Promote Vitality of Capital Market? Analysis from the Perspective of Stock Liquidity". In: *Frontiers in Environmental Science* 11, p. 1132845. DOI: 10.3389/fenvs.2023.1132845.

Choi, Darwin, Zhenyu Gao, and Wenxi Jiang (2020). "Attention to Global Warming". In: *The Review of Financial Studies* 33(3), pp. 1112–1145. DOI: 10.1093/rfs/hhz086.

Christensen, Dane M., George Serafeim, and Anna Sikochi (2022). "Why Is Corporate Virtue in the Eye of the Beholder? The Case of ESG Ratings". In: *The Accounting Review* 97(1), pp. 147–175. DOI: 10.2308/TAR-2019-0506.

Claus, James and Jacob Thomas (2001). "Equity Premia as Low as Three Percent? Evidence from Analysts' Earnings Forecasts for Domestic and International Stock Markets". In: *The Journal of Finance* 56(5), pp. 1629–1666. DOI: 10.1111/0022-1082.00384.

Dobrick, David, Christian Klein, and Bernhard Zwergel (2025). "ESG as a Risk Factor: A Cross-Market Evaluation". In: *Journal of Sustainable Finance & Investment*.

Easton, Peter D. (2004). "PE Ratios, PEG Ratios, and Estimating the Implied Expected Rate of Return on Equity Capital". In: *The Accounting Review* 79(1), pp. 73–95. DOI: 10.2308/accr.2004.79.1.73.

Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel (2020). "Hedging Climate Change News". In: *The Review of Financial Studies* 33(3), pp. 1184–1216. DOI: 10.1093/rfs/hhz072.

Eskildsen, Marc, Markus Ibert, Theis Ingerslev Jensen, and Lasse Heje Pedersen (2024). "In Search of the True Greenium". Working paper. DOI: 10.2139/ssrn.4744608.

Fama, Eugene F. and James D. MacBeth (1973). "Risk, Return, and Equilibrium: Empirical Tests". In: *Journal of Political Economy* 81(3), pp. 607–636. DOI: 10.1086/260061.

Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan (2001). "Toward an Implied Cost of Capital". In: *Journal of Accounting Research* 39(1), pp. 135–176. DOI: 10.1111/1475-679X.00007.

Gibson, Rajna, Philipp Krueger, and Peter Steffen Schmidt (2021). "ESG Rating Disagreement and Stock Returns". In: *Financial Analysts Journal* 77(4). Earlier version circulated as ECGI Finance Working Paper No. 712/2020, pp. 104–127. DOI: 10.1080/0015198X.2021.1963186.

Gormsen, Niels Joachim, Kilian Huber, and Sangmin Simon Oh (2024). *Climate Capitalists*. NBER Working Paper 32933. National Bureau of Economic Research. DOI: 10.3386/w32933.

Hara, Kazuhiko and Thorsten Hens (2025). "Shareholder Engagement and the ESG-CAPM". In: *Working Paper*.

Hartzmark, Samuel M. and Abigail B. Sussman (2019). "Do Investors Value Sustainability? A Natural Experiment Examining Ranking and Fund Flows". In: *The Journal of Finance* 74(6), pp. 2789–2837. DOI: 10.1111/jofi.12841.

Hens, Thorsten and Moritz Trutwin (2024). "An ESG-modified Capital Asset Pricing Model". In: *Swiss Finance Institute Research Paper* (24-01).

Jensen, Scott, Bryan T. Kelly, and Lasse Heje Pedersen (2023). *Global Stock Returns and Characteristics Dataset*.

Johnstone, Louis and Jeremy Grant (2023). "Zero-Beta ESG Projects and the Cost of Capital". In: *Finance Research Letters* 51, p. 102346.

Karolyi, George Andrew, Ying Wu, and William W. Xiong (2025). "Understanding Green Returns in Global Equity Markets: The Role of Representation Bias". In: *Working Paper*.

Krueger, Philipp, Zacharias Sautner, and Laura T. Starks (2020). "The Importance of Climate Risks for Institutional Investors". In: *The Review of Financial Studies* 33(3), pp. 1067–1111. DOI: 10.1093/rfs/hhz137.

Li, Kai, Po-Hsuan Hsu, and Chi-Yang Tsou (2023). "The Pollution Premium". In: *The Journal of Finance* 78(3), pp. 1343–1392. DOI: 10.1111/jofi.13217.

Luo, Di (2022). "ESG, Liquidity, and Stock Returns". In: *Journal of International Financial Markets, Institutions and Money* 78, p. 101526. DOI: 10.1016/j.intfin.2022.101526.

Nakagawa, Kei, Keisuke Morita, and Ryuta Sakemoto (2025). "Stochastic ESG Scores and Nonpecuniary ESG Preferences: An Extension to CAPM". In: *Finance Research Letters* 79, p. 107179. DOI: 10.1016/j.frl.2025.107179.

Oestreich, A. Marcel and Ilias Tsiakas (2015). "Carbon Emissions and Stock Returns: Evidence from the EU Emissions Trading Scheme". In: *Journal of Banking & Finance* 58, pp. 294–308. DOI: 10.1016/j.jbankfin.2015.05.005.

Ohlson, James A. and Beate E. Jüttner-Nauroth (2005). "Expected EPS and EPS Growth as Determinants of Value". In: *Review of Accounting Studies* 10(2–3), pp. 349–365. DOI: 10.1007/s11142-005-1535-3.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor (2021). "Sustainable Investing in Equilibrium". In: *Journal of Financial Economics* 142(2), pp. 550–571. DOI: 10.1016/j.jfineco.2020.12.011.

— (2022). "Dissecting Green Returns". In: *Journal of Financial Economics* 146(2), pp. 403–424. DOI: 10.1016/j.jfineco.2022.07.007.

— (2025). "Sustainable Investing". In: *Annual Review of Financial Economics* 17(Volume 17, 2025), pp. 321–341. DOI: https://doi.org/10.1146/annurev-financial-113023-111705.

Pedersen, Lasse Heje, Sean Fitzgibbons, and Lukasz Pomorski (2021). "Responsible Investing: The ESG-Efficient Frontier". In: *Journal of Financial Economics* 142(2), pp. 572–597. DOI: 10.1016/j.jfineco.2020.11.001.

Rudkin, Mark, Bruce Robertson, and Yi Zhang (2023). "Information Content of Sustainability Index Recomposition: Evidence from the Dow Jones Sustainability Index". In: *Working Paper*.

Shanken, Jay (1992). "On the Estimation of Beta-Pricing Models". In: *The Review of Financial Studies* 5(1), pp. 1–33. DOI: 10.1093/rfs/5.1.1.

Zerbib, Olivier David (2022). "The Sustainable Capital Asset Pricing Model (S-CAPM)". In: *Working Paper*.

8 Tables and Figures

Table 1: Fama–MacBeth Regressions of Implied Cost of Capital (ICC)

	Spec A	Spec B	Spec C	Spec D
Intercept	0.0855***	0.0668***	0.0816***	0.0663***
	(71.75)	(46.99)	(79.88)	(32.74)
Green beta	-0.0012***	-0.0018***	-0.0013**	-0.0012**
	(-2.76)	(-2.64)	(-2.16)	(-2.39)
Market beta		0.0163***	0.0159***	0.0132***
		(18.37)	(18.47)	(19.13)
ESG score level			-0.0032***	-0.0023***
			(-8.77)	(-8.88)
log(ME)				0.0025***
				(9.39)
log(B/M)				0.0065***
				(16.32)
Short-term reversal (1–0)				-0.0221***
				(-8.62)
Momentum (12–1)				-0.0130***
				(-7.69)
Long-term reversal (60–12)				-0.0051***
				(-7.27)
R^2 (avg.)	0.0020	0.1006	0.1101	0.2278
Months	152	152	152	152
Firm-month observations	144,275	133,969	133,969	126,520

Notes. This table reports time-series averages of monthly Fama–MacBeth cross-sectional regressions of the implied cost of capital (ICC) on green beta and firm characteristics. The dependent variable is $\mathrm{ICC}_{i,t}$ in annual decimal units, constructed as the equal-weighted average of model-based ICC estimates and winsorized within month at the 1st/99th percentiles. Green beta $\widehat{\beta}^{g}_{i,t}$ is estimated from the rolling 36-month ESG–ESG regression in Eq. (5.2) using the value-weighted market ESG score; $\widehat{\beta}^{g}_{i,t}$ and the ESG score level are winsorized within month at the 1st/99th percentiles, demeaned each month, and scaled by fixed cross-sectional standard deviations (time-series averages of within-month dispersions), so the reported coefficients can be interpreted as changes in annual ICC per one-standard-deviation difference on a constant scale. Each monthly cross-section is estimated by OLS with equal firm weights. Reported t-statistics (in parentheses) are based on Newey–West HAC standard errors with six lags applied to the time series of monthly slope estimates. Empty cells indicate that the regressor is not included in the specification. Coefficients may be multiplied by 10,000 to express effects in annual basis points. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table 2: State-Dependent Pricing of Green Beta and ESG Score: Climate-Concern Shocks (MCCC)

	Price of green beta $\lambda_{g,t}$	Price of ESG score $\pi_{s,t}$
Intercept	-11.99***	-22.90***
	(-3.30)	(-12.33)
Shock_{t-1}	7.39***	-2.67*
	(2.87)	(-1.80)
Months	152	152

Notes. Each month t, we estimate the ICC cross-section in Eq. (5.3) using the full control set and obtain the monthly prices of green-beta exposure, $\widehat{\lambda}_{g,t}$, and ESG levels, $\widehat{\pi}_{s,t}$. Green beta and the ESG score level are demeaned each month and expressed per one-standard-deviation difference using the fixed (time-invariant) scaling defined in Section 5 (time-series averages of within-month cross-sectional dispersions). We report coefficients in *annual basis points per a one-standard-deviation difference in exposure/level* (i.e., $10{,}000 \times \widehat{\lambda}_{g,t}$ and $10{,}000 \times \widehat{\pi}_{s,t}$). Shock_{t-1} is the standardized climate-concern shock (lagged by one month). Because the shock is standardized, the Shock_{t-1} coefficient measures the change in the annual ICC price (bps per 1σ exposure/level) associated with a $+1\sigma$ climate-concern shock. t-statistics (in parentheses) use Newey–West HAC standard errors with six lags. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table 3: Time Trend in the Green-Beta Premium and Taste (Score) Channel

	(1) $g_{\beta,t}$	(2) $g_{s,t}$	(3) Prem_{Q5-Q1}	(4) $\Delta\beta_t^g$
Intercept	-47.22***	0.08	-26.88***	1.84***
	(-4.34)	(0.03)	(-3.57)	(3.69)
Trend[0,1]	70.33***	-45.88***	57.26***	3.16***
	(4.56)	(-10.39)	(4.42)	(3.73)
Months	152	152	152	152

Notes. This table regresses monthly pricing objects on a linear time trend normalized to run from 0 at the start of the estimation sample to 1 at the end. Columns (1) and (2) report the annual basis-point prices of standardized green beta and standardized ESG level, respectively (i.e., $10{,}000 \times \widehat{\lambda}_{g,t}$ and $10{,}000 \times \widehat{\pi}_{s,t}$). Column (3) reports the implied annual ICC premium between the highest and lowest green-beta quintiles, $Prem_{Q5-Q1,t} = 10{,}000 \times \widehat{\lambda}_{g,t} \times \Delta\beta_t^g$, where $\Delta\beta_t^g$ is the within-month Q5–Q1 spread in green beta exposures (column (4), in beta units). Reported t-statistics are based on Newey–West HAC standard errors with six lags. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table 4: Industry portfolios: Fama–MacBeth pricing of green beta in ICC cross-sections

	Spec A	Spec B	Spec C	Spec D	A+logME	B+logBM
Industry green beta	-0.0033***	-0.0026***	-0.0018**	-0.0007	-0.0022**	-0.0021**
	(-3.42)	(-2.97)	(-2.55)	(-1.22)	(-2.55)	(-2.59)
Industry market beta		0.0188***	0.0157***	0.0154***	0.0213***	0.0162***
		(16.17)	(11.76)	(11.66)	(17.39)	(14.63)
Industry ESG level			-0.0053***	-0.0036***		
			(-4.97)	(-3.94)		
log(ME)					0.0028***	
					(5.77)	
log(B/M)				0.0076***		0.0080***
				(8.08)		(13.84)
Short-term reversal				-0.0321***		
(1–0)				(-3.88)		
Momentum				-0.0019		
(12–1)				(-0.45)		
Long-term reversal				0.0007		
(60–12)				(0.34)		
Constant	0.0838***	0.0635***	0.0923***	0.0945***	0.0266***	0.0762***
	(52.34)	(40.67)	(15.54)	(19.31)	(4.48)	(46.58)
R^2 (avg.)	0.0598	0.2449	0.2942	0.4896	0.3259	0.3455
Months	152	152	152	152	152	152

Notes. This table reports time-series averages of monthly Fama–MacBeth cross-sectional regressions run across Fama–French 49 industries. The dependent variable is the value-weighted industry ICC (annual decimal units), computed as the market-cap-weighted average of firm ICC within each industry-month. Industry green beta is estimated from the rolling 36-month regression of the industry ESG score on the value-weighted market ESG score *excluding* that industry (leave-one-industry-out), as described in Section 6.4, and is winsorized within month at the 1st/99th percentiles. Each month, the cross-section is estimated by OLS with equal industry weights; reported t-statistics (in parentheses) are based on Newey–West HAC standard errors (6) applied to the time series of monthly slopes. Empty cells indicate a regressor is not included in the specification. Coefficients may be multiplied by 10,000 to express effects in annual basis points. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table 5: Within- vs between-industry decomposition of green-beta pricing

	EW industry mean	VW industry mean
Within-industry component	-0.0011***	-0.0010***
	(-3.63)	(-2.69)
Between-industry mean	-0.0004	-0.0010
	(-0.65)	(-0.91)
Market beta	0.0132***	0.0132***
	(25.29)	(18.63)
ESG score level	-0.0022***	-0.0021***
	(-12.00)	(-8.53)
log(ME)	0.0025***	0.0025***
	(12.54)	(9.05)
log(B/M)	0.0065***	0.0065***
	(22.24)	(17.05)
Short-term reversal (1–0)	-0.0221***	-0.0222***
	(-9.09)	(-8.50)
Momentum (12–1)	-0.0131***	-0.0131***
	(-10.16)	(-7.73)
Long-term reversal (60–12)	-0.0051***	-0.0051***
	(-10.16)	(-7.24)
Constant	0.0661***	0.0659***
	(44.63)	(34.81)
R^2 (avg.)	0.2307	0.2329
Months	152	152

Notes. This table decomposes firm green beta into within- and between-industry components and reports Fama–MacBeth ICC cross-sectional regressions that include both components simultaneously. In each month t, we write $\beta^g_{i,t} = (\beta^g_{i,t} - \bar{\beta}^g_{k(i),t}) + \bar{\beta}^g_{k(i),t}$, where $\bar{\beta}^g_{k,t}$ is the industry mean computed either equal-weighted across firms (column 1) or value-weighted by market equity (column 2). The dependent variable is $\mathrm{ICC}_{i,t}$ in annual decimal units. All regressions include the full set of controls shown in the table; the greenness regressors follow the paper's fixed-scale standardization so coefficients can be interpreted per one-standard-deviation difference on a constant scale. Reported t-statistics are based on Newey–West standard errors (6) applied to the time series of monthly slopes. Coefficients may be multiplied by 10,000 to express effects in annual basis points. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table 6: Within-industry neutral green-beta portfolios

	VW / equal	EW / equal	EW / cap
Panel A: One-month-ahead excess returns			
Q1 (Low)	0.0123	0.0118	0.0112
	(3.42)	(2.80)	(2.72)
Q2	0.0119	0.0126	0.0120
	(3.53)	(3.06)	(2.87)
Q3	0.0107	0.0098	0.0096
	(3.60)	(2.42)	(2.20)
Q4	0.0099	0.0114	0.0106
	(3.03)	(3.02)	(2.75)
Q5 (High)	0.0098	0.0100	0.0094
	(3.09)	(2.65)	(2.48)
H–L (Q5–Q1)	-0.0025	-0.0018	-0.0018
	(-1.73)	(-1.29)	(-1.35)
Months	145	145	145
Panel B: Implied cost of capital (ICC)			
Q1 (Low)	0.0812	0.0844	0.0844
	(51.80)	(62.43)	(67.06)
Q2	0.0807	0.0843	0.0844
	(44.51)	(61.95)	(63.35)
Q3	0.0798	0.0837	0.0834
	(66.03)	(56.57)	(58.26)
Q4	0.0803	0.0840	0.0839
	(56.17)	(59.24)	(59.23)
Q5 (High)	0.0794	0.0810	0.0814
	(52.06)	(70.63)	(75.03)
H–L (Q5–Q1)	-0.0017	-0.0034	-0.0030
	(-1.72)	(-4.28)	(-3.87)
Months	151	151	151

Notes. Each month, firms are sorted *within* Fama–French 49 industries into green-beta quintiles using $\widehat{\beta}^{g}_{i,t}$ (industry-demeaned sorting implied by the within-industry construction). Column labels indicate the weighting scheme used to form and aggregate portfolios: VW (*value-weighted*) and EW (*equal-weighted*) refer to firm weights *within* each industry; *equal* and *cap* refer to how industry-level portfolios are aggregated *across* industries (equal-weighting industries vs. weighting industries by total industry market capitalization). We require at least 35 firms per industry-month with non-missing $\hat{\beta}^{g}_{i,t}$ (and the relevant outcome) to form quintiles; months failing this requirement are dropped. Panel A reports subsequent-month excess returns (monthly decimals). Panel B reports implied cost of capital (annual decimals). H–L denotes Q5–Q1. t-statistics in parentheses are based on Newey–West standard errors (6).

Table 7: Within-industry double sorts on green beta and ESG level: implied cost of capital

	ESG Q1	ESG Q2	ESG Q3	ESG Q4	ESG Q5	H–L (ESG5–ESG1)
GB Q1 (Low)	0.0892	0.0796	0.0783	0.0815	0.0758	-0.0134
						(-3.57)
GB Q2	0.0876	0.0832	0.0812	0.0808	0.0753	-0.0121
						(-4.72)
GB Q3	0.0882	0.0793	0.0805	0.0817	0.0727	-0.0155
						(-5.87)
GB Q4	0.0798	0.0768	0.0763	0.0791	0.0758	-0.0040
						(-2.09)
GB Q5 (High)	0.0774	0.0743	0.0790	0.0802	0.0742	-0.0031
						(-1.25)
H–L (GB5–GB1)	-0.0127	-0.0049	0.0013	-0.0014	-0.0006	0.0121
	(-2.73)	(-2.24)	(0.54)	(-0.58)	(-0.45)	(2.81)

Notes. This table reports time-series averages of value-weighted portfolio implied costs of capital (ICC, annual decimals) for 5×5 within-industry double sorts on green beta and ESG level. In each month t and within each Fama–French 49 industry, firms are sorted into ESG quintiles using the within-industry ESG measure and then, within each ESG quintile, into green-beta quintiles using the within-industry green-beta measure (25 portfolios). ICC is measured in month t and portfolios are value-weighted using lagged market equity. The sample requires at least 15 firms per industry and per cell-month. H–L spreads are Q5–Q1 differences; Newey–West (6) t-statistics for H–L spreads are in parentheses.

Table 8: Spanning regressions for the Green-Beta factor (industry-neutral, spread aggregation)

	Dependent variable: GB_t				
	(1) α	(2) FF3	(3) FF3+Mom	(4) FF5	(5) FF5+Mom
MKT-RF		-0.009	-0.011	-0.012	-0.013
		(-0.22)	(-0.27)	(-0.29)	(-0.31)
SMB		-0.061	-0.063	-0.055	-0.057
		(-1.05)	(-1.11)	(-0.83)	(-0.88)
HML		0.022	0.020	0.034	0.032
		(0.85)	(0.69)	(0.77)	(0.73)
RMW				0.026	0.024
				(0.37)	(0.33)
CMA				-0.035	-0.032
				(-0.47)	(-0.44)
Mom			-0.012		-0.007
			(-0.29)		(-0.16)
α (bps/month)	-24.8**	-23.9**	-23.5**	-24.0**	-23.7**
	(-2.15)	(-2.33)	(-2.31)	(-2.40)	(-2.37)
Observations	152	152	152	152	152
R^2	0.000	0.010	0.011	0.012	0.012

Notes. The dependent variable is GB_t, the monthly Green-Beta factor return constructed using the industry-neutral *spread-aggregation* procedure described in Section 6.5: within each Fama–French 49 industry, firms are sorted into green-beta quintiles using $\widehat{\beta}^{g}_{i,t}$ and the within-industry return spread (Q5–Q1) is computed; these industry spreads are then averaged across industries to hold industry weights fixed by construction. The table reports spanning regressions of GB_t on Fama–French factors (FF3 and FF5) and momentum. GB_t and the factors are monthly excess returns in decimals; α is reported in basis points per month. t-statistics (in parentheses) are based on Newey–West standard errors with six lags.

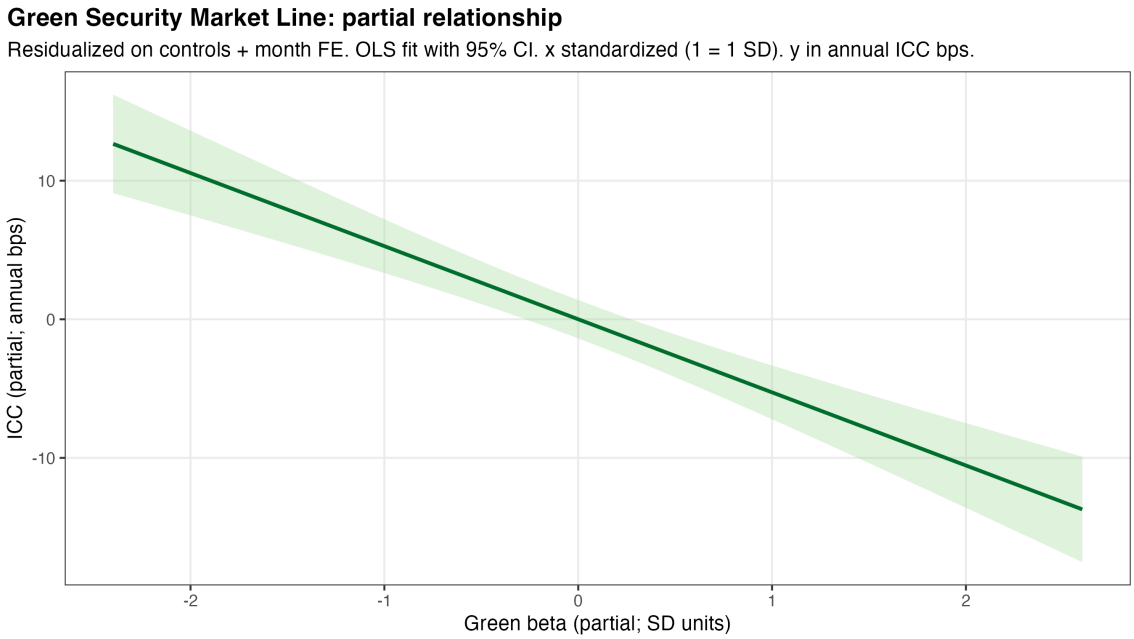


Figure 1: Green Security Market Line: partial relationship. The figure plots the fitted OLS relation between implied cost of capital (ICC) and green beta after residualizing both variables on the full control set and month fixed effects. Green beta is standardised so that one unit corresponds to one cross-sectional standard deviation. The y-axis reports the partial component of ICC in annual basis points. The shaded band is a 95% confidence interval.

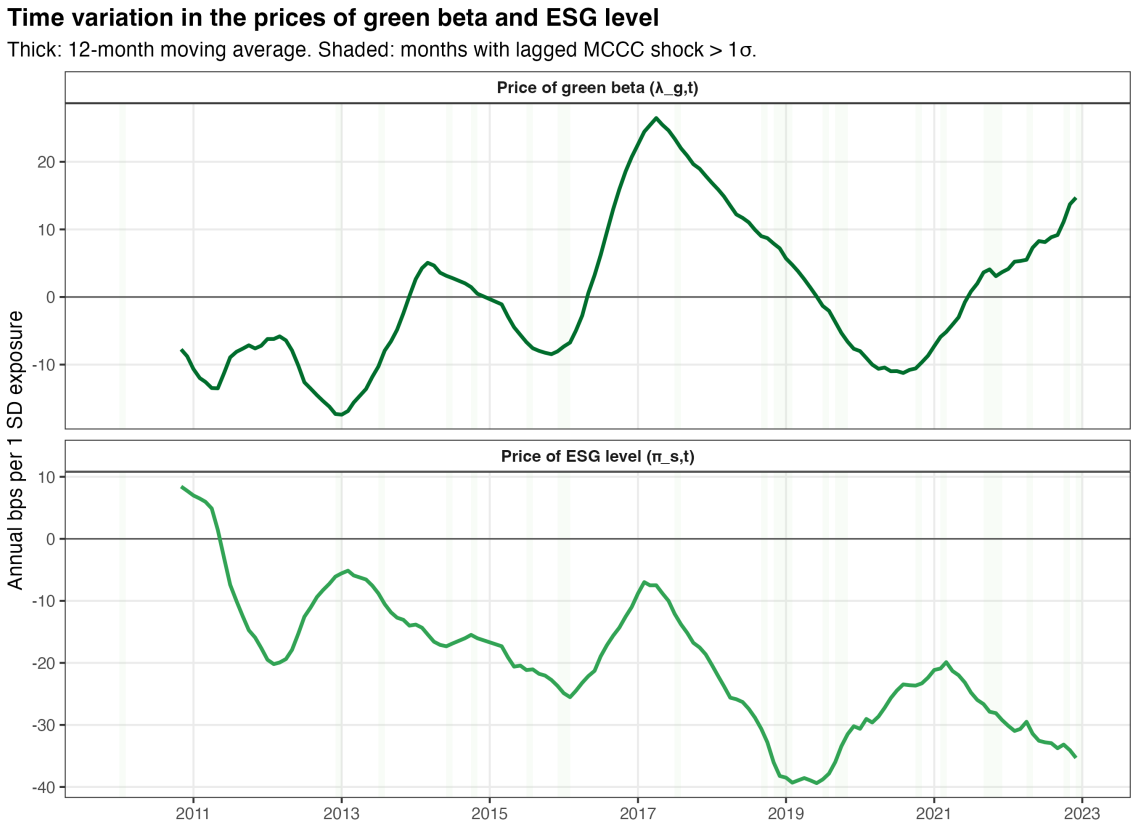


Figure 2: Time variation in the prices of green-beta exposure and ESG-score level. Each panel plots the 12-month moving average of the monthly price of green beta, λ_t^{β}, and the price of ESG-score level, π_t^s, estimated from month-by-month ICC cross-sectional regressions. Prices are expressed in annual basis points per one-standard-deviation difference in the relevant characteristic. Shaded regions denote climate-concern months, defined as months in which the lagged standardised MCCC innovation exceeds 1σ.

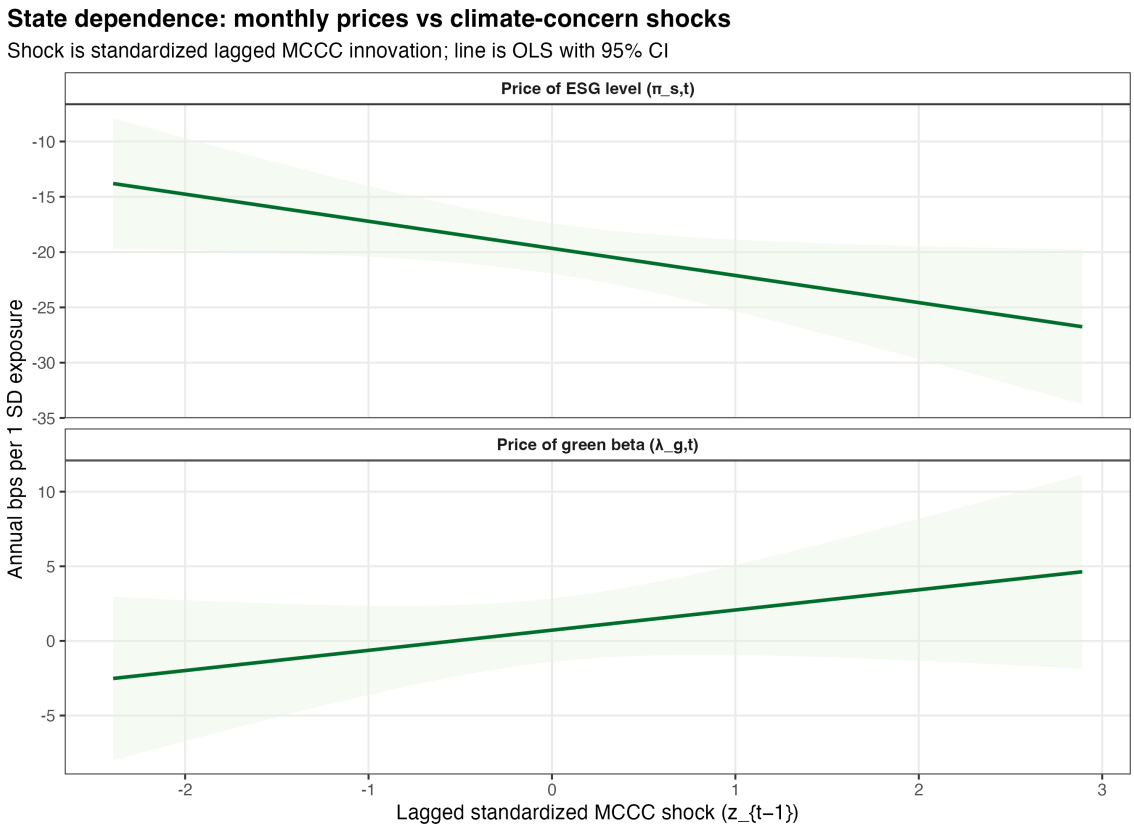


Figure 3: State dependence of monthly prices on climate-concern shocks. The figure relates the monthly price series to the lagged standardised MCCC innovation. The solid line is the fitted OLS relation with a 95% confidence band. Prices are expressed in annual basis points per one-standard-deviation exposure.

A Appendix

This appendix collects supplementary analyses that support and extend the main results. Appendix A.1 provides the summary statistics. Appendix A.2 reports realized-return counterparts to the baseline ICC evidence and additional portfolio diagnostics. Appendix A.3 provides complementary time-variation evidence via a pre-/post-2016 split. Appendix A.4 reports robustness checks for the climate-salience state-dependence results, including alternative shock timing and nonlinear specifications. Appendix A.5 repeats the main two-channel tests using MSCI's Environmental Pillar Score and Trucost Scope 1+ Scope 2 Carbon Emissions Intensity. Appendix A.6 collects figures that visually summarize the market ESG benchmark, the distribution of estimated green betas, and time-series diagnostics related to state dependence.

A.1 Summary Statistics

This subsection reports descriptive statistics for the main firm-level variables used in the ICC cross-sectional regressions, and summarizes the key pricing objects used throughout the paper. Table A.1 reports summary statistics and Table A.2 reports pairwise correlations. Table A.3 reports average ICC across green-beta-sorted portfolios.

Table A.1: Summary statistics (core variables)

Variable	N	Mean	SD	P10	P50	P90
ICC ($ICC_{i,t}$)	126,520	0.083	0.032	0.051	0.079	0.118
Green beta (β_i^g)	126,520	0.680	1.638	-0.901	0.494	2.613
ESG score level ($es_{i,t}$)	126,520	4.701	0.984	3.500	4.700	6.000
Market beta (β_i^m)	126,520	1.152	0.558	0.488	1.111	1.825
log(ME)	126,520	8.893	1.408	7.151	8.752	10.810
log(B/M)	126,520	-0.980	0.919	-2.099	-0.875	0.013
Momentum (12–1)	126,520	0.085	0.319	-0.288	0.098	0.433
Long-term reversal (60–12)	126,520	0.364	0.574	-0.310	0.398	1.001

Notes. This table reports summary statistics for the firm-month sample used in the baseline ICC cross-sectional regressions. $ICC_{i,t}$ is the implied cost of capital in annual decimal units, constructed as the equal-weighted average of model-based ICC estimates and winsorized within month at the 1st/99th percentiles. Green beta is the rolling 36-month slope estimate $\widehat{\beta}_{i,t}^g$ from Eq. (5.2) using the value-weighted market ESG score. The ESG score level is MSCI's raw weighted-average ESG score (0–10). All variables are shown in raw units; in the cross-sectional pricing tests, green beta and the ESG level are demeaned each month and scaled by fixed cross-sectional standard deviations (time-series averages of within-month dispersions) as described in Section 5.

Table A.2: Correlation matrix (core variables)

	ICC	Green beta	ESG Score	Market beta	log(ME)	log(B/M)	Mom (12–1)	Long-Term (60–12)
ICC	1.00							
Green beta	-0.01	1.00						
ESG Score	-0.08	-0.02	1.00					
Market beta	0.28	-0.01	-0.07	1.00				
log(ME)	-0.05	-0.01	0.17	-0.21	1.00			
log(B/M)	0.32	-0.03	-0.14	0.13	-0.29	1.00		
Momentum (12–1)	-0.14	-0.02	0.02	-0.01	0.18	-0.25	1.00	
Long-term reversal (60–12)	-0.26	0.06	0.02	-0.21	0.19	-0.37	-0.11	1.00

Notes. This table reports pairwise Pearson correlations computed over the firm-month panel in Table A.1. Variables are measured contemporaneously (month t). The table shows the lower triangle of the correlation matrix.

Table A.3: ICC by green-beta quintiles (time-series averages; Newey–West t-statistics)

Portfolio	Mean ICC	t	Months
Q1	0.0827	48.10	152
Q2	0.0851	49.61	152
Q3	0.0857	44.02	152
Q4	0.0838	54.35	152
Q5	0.0824	38.92	152
Q5–Q1	-0.0003	-0.29	152

Notes. Each month, firms are sorted into quintiles based on green beta $\widehat{\beta}^{g}_{i,t}$. The table reports time-series averages of value-weighted portfolio ICC (annual decimals). Q5–Q1 denotes the difference between the highest- and lowest-green-beta portfolios. Reported t-statistics for the portfolio means are based on Newey–West standard errors for the time series of monthly portfolio averages. Unconditional green-beta sorts are not strongly monotone in ICC, consistent with the near-zero unconditional correlation (Table A.2); the pricing relation emerges in multivariate specifications and within-industry designs.

A.2 Additional Evidence Using Realized Returns

This subsection reports return-based counterparts to the ICC evidence in Table 1 and to the portfolio evidence in Table 8. Because one-month-ahead realized returns are substantially noisier than implied expected returns, these tests are designed as diagnostic complements rather than as the primary empirical design. Table A.4 replicates the baseline Fama–MacBeth cross-sectional specifications using subsequent-month excess returns ($r^{e}_{i,t+1}$). As expected given the low signal-to-noise ratio of monthly realized returns, the estimated greenness prices are much less precise than their ICC counterparts: the green-beta coefficient remains negative across specifications but is not statistically different from zero, while the ESG-level coefficient is economically close to zero. Table A.5 repeats the within-industry double-sort design of Table 8 using subsequent-month excess returns and shows that the return ordering is directionally consistent with the ICC spreads, albeit with larger sampling variation. Finally, Tables A.6–A.7 complement the spanning regressions in Table 4 by considering alter-

native constructions of the green-beta factor (GB_t), including an unconditional sort and an industry-neutral construction with firm-level aggregation. Taken together, these return-based diagnostics support the interpretation that the green-beta channel captures distinct variation in expected returns, while also illustrating why ICC-based evidence is sharper at the monthly horizon.

Table A.4: Fama–MacBeth Regressions Using One-Month-Ahead Excess Returns

	Spec A	Spec B	Spec C	Spec D
Intercept	0.0107***	0.0086***	0.0080**	0.0124*
	(3.11)	(4.03)	(2.15)	(1.88)
Green beta	-0.0008	-0.0008	-0.0010	-0.0009
	(-0.87)	(-0.98)	(-1.10)	(-1.03)
Market beta		0.0016	0.0016	0.0003
		(0.66)	(0.66)	(0.12)
ESG score level			0.0002	0.0002
			(0.35)	(0.48)
log(ME)				-0.0005
				(-1.15)
log(B/M)				-0.0006
				(-0.96)
Short-term reversal (1–0)				-0.0119*
				(-1.71)
Momentum (12–1)				0.0036
				(1.08)
Long-term reversal (60–12)				0.0004
				(0.23)
R^2 (avg.)	0.0034	0.0462	0.0508	0.1139

Notes. This table reports time-series averages of monthly Fama–MacBeth cross-sectional regressions where the dependent variable is the subsequent-month excess stock return, $r^e_{i,t+1}$ (monthly decimals). Regressors are measured at month t and follow the same definitions as in Table 1. Green beta and ESG level are demeaned each month and scaled by fixed cross-sectional standard deviations (time-series averages of within-month dispersions), so coefficients can be interpreted as differences in next-month expected returns per one-standard-deviation difference on a constant scale. Reported t-statistics (in parentheses) are based on Newey–West (6) standard errors applied to the time series of monthly slope estimates; for the coefficient on the generated green-beta regressor we additionally apply the Shanken (1992) adjustment. The implied Shanken inflation factor is approximately one in this sample, so reported t-statistics are unchanged at the displayed precision. Empty cells indicate that the regressor is not included in the specification. Coefficients may be multiplied by 10,000 to express effects in monthly basis points. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table A.5: Within-industry double sorts on green beta and ESG level: subsequent-month excess returns

	ESG Q1	ESG Q2	ESG Q3	ESG Q4	ESG Q5	H–L (ESG5–ESG1)
GB Q1 (Low)	0.0186	0.0077	0.0091	0.0094	0.0108	-0.0078
						(-1.36)
GB Q2	0.0137	0.0106	0.0126	0.0081	0.0086	-0.0043
						(-0.87)
GB Q3	0.0041	0.0116	0.0097	0.0088	0.0108	0.0065
						(1.63)
GB Q4	0.0087	0.0125	0.0094	0.0113	0.0104	0.0015
						(0.47)
GB Q5 (High)	0.0063	0.0058	0.0071	0.0086	0.0099	0.0034
						(1.09)
H–L (GB5–GB1)	-0.0125	-0.0022	-0.0027	-0.0025	-0.0017	0.0108
	(-1.98)	(-0.48)	(-0.60)	(-0.77)	(-0.44)	(1.41)

Notes. This table reports time-series averages of subsequent-month excess returns (monthly decimals) for the same 5×5 within-industry double sorts as in Table 7. In each month and within each Fama–French 49 industry, firms are first sorted into ESG quintiles and then, within each ESG quintile, into green-beta quintiles. Portfolio returns are value-weighted using market equity at portfolio formation. H–L spreads are Q5–Q1 differences; Newey–West (6) t-statistics for H–L spreads are in parentheses.

Table A.6: Spanning regressions for the Green-Beta factor (unconditional sort)

	Dependent variable: GB_t				
	(1) α	(2) FF3	(3) FF3+Mom	(4) FF5	(5) FF5+Mom
MKT-RF		-0.164	-0.211	-0.193	-0.225
		(-0.92)	(-0.91)	(-0.93)	(-0.92)
SMB		-0.057	-0.110	-0.051	-0.099
		(-0.66)	(-1.11)	(-0.47)	(-1.11)
HML		-0.118	-0.175	0.016	-0.079
		(-0.95)	(-0.88)	(0.33)	(-0.56)
RMW				0.089	0.062
				(0.34)	(0.27)
CMA				-0.336	-0.226
				(-0.99)	(-1.08)
Mom			-0.230		-0.202
			(-0.74)		(-0.70)
α (bps/month)	-86.4	-69.8	-60.2	-62.0	-56.2
	(-1.37)	(-1.47)	(-1.56)	(-1.39)	(-1.40)
Observations	152	152	152	152	152
R^2	0.000	0.015	0.020	0.015	0.020

Notes. GB_t is a high–minus–low green-beta factor constructed each month using an unconditional (full-universe) green-beta quintile sort; the factor return is the Q5–Q1 return spread. The table reports spanning regressions of GB_t on Fama–French factors and momentum. GB_t and the factors are monthly excess returns in decimals; α is reported in basis points per month. Newey–West (6) t-statistics are in parentheses.

Table A.7: Spanning regressions for the Green-Beta factor (industry-neutral, firm aggregation)

	Dependent variable: GB_t				
	(1) α	(2) FF3	(3) FF3+Mom	(4) FF5	(5) FF5+Mom
MKT-RF		-0.255	-0.299	-0.290	-0.318
		(-1.40)	(-1.28)	(-1.38)	(-1.29)
SMB		0.006	-0.044	0.028	-0.014
		(0.06)	(-0.42)	(0.21)	(-0.12)
HML		-0.141	-0.194	0.0001	-0.083
		(-1.12)	(-0.98)	(0.00)	(-0.58)
RMW				0.143	0.119
				(0.55)	(0.51)
CMA				-0.369	-0.273
				(-1.06)	(-1.23)
Mom			-0.213		-0.177
			(-0.70)		(-0.63)
α (bps/month)	-87.8	-62.1	-53.2	-54.5	-49.4
	(-1.43)	(-1.37)	(-1.48)	(-1.29)	(-1.32)
Observations	152	152	152	152	152
R^2	0.000	0.019	0.023	0.021	0.024

Notes. GB_t is an industry-neutral high–minus–low green-beta factor constructed by sorting firms into green-beta quintiles *within* each Fama–French 49 industry and aggregating the long and short legs across industries at the firm level (firm aggregation) to hold industry weights fixed in both legs. The table reports spanning regressions of GB_t on Fama–French factors and momentum. GB_t and the factors are monthly excess returns in decimals; α is reported in basis points per month. Newey–West (6) t-statistics are in parentheses.

A.3 Additional Time-Variation Evidence

This subsection provides complementary time-variation diagnostics for the paper's two pricing channels: the exposure channel (green-beta pricing) and the level channel (ESG-score pricing). Table A.8 summarizes the evolution of the monthly pricing objects using a transparent pre-/post-2016 split, which serves as a change-in-means counterpart to the normalized-trend regressions in Table 3. The evidence reinforces that the exposure and level channels exhibit distinct time-series behavior, consistent with the two-channel interpretation emphasized in the main text.

Table A.8: Pre-2016 vs Post-2016 Comparison of Green-Beta Pricing and Taste Channel

	(1) $\lambda_{g,t}$	(2) Prem_{Q5-Q1}	(3) $\Delta\beta_t^g$	(4) $\pi_{s,t}$
Pre-2016 mean (Intercept)	-27.12***	-21.95***	1.88***	-10.13***
	(-3.73)	(-4.28)	(5.76)	(-4.23)
Post-2016 mean (Pre + Shift)	3.79	20.92	4.67	-31.48
Post-2016 shift (Post – Pre)	30.91***	42.87***	2.79***	-21.35***
	(4.13)	(4.78)	(5.11)	(-6.00)
Months	152	152	152	152

Notes. This table reports a pre-/post-2016 split for monthly pricing objects. The pre-2016 mean is the time-series average prior to January 2016 in the estimation sample; the post-2016 shift is the difference between the post-2016 mean and the pre-2016 mean. The row "Post-2016 mean (Pre + Shift)" reports the implied post-2016 mean (intercept plus shift). Columns (1), (2), and (4) are in annual basis points; column (3) is in green-beta units. Reported t-statistics (in parentheses) are based on Newey–West standard errors (6) applied to the monthly series. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

A.4 Additional State-Dependence Specifications

This subsection reports robustness checks for the state-dependence results in Table 2. Table A.9 allows both contemporaneous and lagged climate-concern shocks (MCCC) to predict the monthly prices, guarding against timing misalignment between investor attention and the pricing objects. Table A.10 applies the same state-dependence design to the implied Q5–Q1 green-beta premium. Tables A.11–A.12 examine nonlinear and pooled alternatives (including regime and equal-weighted-month specifications). Across these specifications, the qualitative pattern is robust: the exposure price responds positively to climate-concern shocks in linear timing variants, while the level price exhibits weaker linear sensitivity but substantially larger shifts in high-concern regimes. This robustness reduces concerns that the baseline state-dependence evidence is driven by a particular functional form or timing convention.

Table A.9: Alternative shock timing for state dependence (2-step FM)

	(1) Shock$_t$	(2) Shock$_{t-1}$	(3) Both
Panel A: $\lambda_{g,t}$ (bps per unit)			
Intercept	-9.91***	-9.98***	-9.85***
	(-3.23)	(-3.30)	(-3.31)
Shock$_t$	3.89**		3.91**
	(2.13)		(2.32)
Shock$_{t-1}$		6.15***	6.16***
		(2.87)	(2.95)
Months	152	152	152
Panel B: $\pi_{s,t}$ (bps per unit)			
Intercept	-22.00***	-21.96***	-22.03***
	(-12.30)	(-12.33)	(-12.52)
Shock$_t$	-2.21		-2.21
	(-1.49)		(-1.52)
Shock$_{t-1}$		-2.56*	-2.56*
		(-1.80)	(-1.84)
Months	152	152	152

Notes. This table reports alternative timing specifications for state-dependent pricing in the second-stage regressions. Monthly price series ($\widehat{\lambda}_{g,t}$ and $\widehat{\pi}_{s,t}$) are obtained from the month-by-month ICC cross-sections (preferred specification) and are expressed in annual basis points (i.e., ICC units multiplied by 10,000). The regressors are standardized MCCC climate-concern shocks entered contemporaneously, lagged, or jointly. Reported t-statistics are based on Newey–West standard errors (6). ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table A.10: State dependence of the implied Q5–Q1 premium

	Prem$_{Q5-Q1}$ (bps)
Intercept	1.86
	(0.33)
Shock$_t$	2.08
	(0.83)
Shock$_{t-1}$	4.88*
	(1.78)
Months	152

Notes. The dependent variable is the implied annual Q5–Q1 green-beta premium in basis points, constructed from the monthly green-beta price and the within-month Q5–Q1 spread in green beta exposures: $Prem_{Q5-Q1,t} = 10{,}000 \times \widehat{\lambda}_{g,t} \times \Delta\beta_t^g$. Regressors are standardized MCCC climate-concern shocks entered contemporaneously and lagged. Reported t-statistics are based on Newey–West standard errors (6). ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Table A.11: Nonlinear state dependence: second-stage (monthly prices)

	(1) Linear	(2) Quadratic	(3) Sign split	(4) High/Low
Panel A: $\lambda_{g,t}$ (bps per unit)				
Intercept	-9.98***	-10.82***	-10.46**	-10.69***
	(-3.30)	(-3.03)	(-2.30)	(-2.89)
Shock_{t-1}	6.15***	5.78**		
	(2.87)	(2.42)		
Shock^2_{t-1}		0.83		
		(0.63)		
Shock^+_{t-1}			6.67**	
			(2.17)	
Shock^-_{t-1}			5.46	
			(0.99)	
High Shock_{t-1}				9.83**
				(2.21)
Low Shock_{t-1}				-6.22
				(-0.76)
Months	152	152	152	152
Panel B: $\pi_{s,t}$ (bps per unit)				
Intercept	-21.96***	-21.66***	-20.54***	-20.55***
	(-12.33)	(-10.71)	(-8.23)	(-10.61)
Shock_{t-1}	-2.56*	-2.43		
	(-1.80)	(-1.65)		
Shock^2_{t-1}		-0.30		
		(-0.27)		
Shock^+_{t-1}			-4.07	
			(-1.56)	
Shock^-_{t-1}			-0.54	
			(-0.18)	
High Shock_{t-1}				-9.60***
				(-2.70)
Low Shock_{t-1}				0.93
				(0.21)
Months	152	152	152	152

Notes. This table reports nonlinear second-stage state-dependence specifications for the monthly price series. Coefficients are reported in annual basis points (ICC units multiplied by 10,000). Column (1) includes the lagged standardized shock linearly; column (2) adds a quadratic term; column (3) splits the lagged shock into positive and non-positive components ($Shock^+ = Shock \cdot \mathbb{1}\{Shock > 0\}$ and $Shock^- = Shock \cdot \mathbb{1}\{Shock \leq 0\}$); and column (4) uses high/low regimes based on the empirical distribution of the standardized shock (e.g., top/bottom tercile). Reported t-statistics are based on Newey–West standard errors with six lags. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

Pooled month-FE versus two-step Fama–MacBeth. Table A.12 estimates state dependence in a single pooled firm–month regression with month fixed effects. Even when observations are weighted by $1/n_t$ so that each month receives equal total weight, the pooled coefficient on $\tilde{\beta}^g_{i,t}$ does not in general equal the simple time-series average of the month-by-month (Fama–MacBeth) slopes. By the Frisch–Waugh–Lovell theorem, the pooled estimate can be written as a weighted average of the month-specific *partial* slopes, with weights proportional to the within-month second moment of the residualized regressor. When the month-specific price of $\tilde{\beta}^g_{i,t}$ and/or the cross-sectional dispersion of $\tilde{\beta}^g_{i,t}$ varies over time, these weights can differ substantially across months, so the pooled estimate may differ from the equal-weighted Fama–MacBeth average (including in sign).

In the sample underlying Column (1), the pooled estimate is 3.55 bps per one-standard-deviation of $\tilde{\beta}^g_{i,t}$; the corresponding variance-weighted average of the month-by-month partial slopes is 4.71 bps, while the simple equal-weighted average is −12.25 bps.

Table A.12: Pooled month fixed effects state dependence (eq-weighted months)

	(1) Linear	(2) Quadratic	(3) Sign split	(4) High/Low
$\tilde{\beta}^g_{i,t}$	3.55***	5.10***	6.68***	4.92***
	(3.98)	(5.05)	(5.34)	(4.92)
$\tilde{s}_{i,t}$	-17.82***	-16.76***	-15.06***	-15.70***
	(-12.51)	(-9.14)	(-5.82)	(-9.72)
$\tilde{\beta}^g_{i,t}\times$ Shock$_{t-1}$	-0.35	0.30		
	(-0.43)	(0.34)		
$\tilde{\beta}^g_{i,t}\times$ Shock$^2_{t-1}$		-1.78***		
		(-3.33)		
$\tilde{s}_{i,t}\times$ Shock$_{t-1}$	-3.51*	-3.13		
	(-1.80)	(-1.56)		
$\tilde{s}_{i,t}\times$ Shock$^2_{t-1}$		-1.02		
		(-0.66)		
$\tilde{\beta}^g_{i,t}\times$ Shock$^+_{t-1}$			-4.10***	
			(-2.96)	
$\tilde{\beta}^g_{i,t}\times$ Shock$^-_{t-1}$			4.40***	
			(2.81)	
$\tilde{s}_{i,t}\times$ Shock$^+_{t-1}$			-6.54*	
			(-1.80)	
$\tilde{s}_{i,t}\times$ Shock$^-_{t-1}$			0.32	
			(0.08)	
$\tilde{\beta}^g_{i,t}\times$ High Shock$_{t-1}$				-6.27**
				(-2.58)
$\tilde{\beta}^g_{i,t}\times$ Low Shock$_{t-1}$				-3.63
				(-1.59)
$\tilde{s}_{i,t}\times$ High Shock$_{t-1}$				-14.00***
				(-2.75)
$\tilde{s}_{i,t}\times$ Low Shock$_{t-1}$				0.31
				(0.06)
Observations	126520	126520	126520	126520
Months (FE)	152	152	152	152
Within R^2	0.188	0.188	0.188	0.188

Notes. This table reports pooled firm–month regressions with month fixed effects that estimate state-dependent pricing in a single step. The dependent variable is $ICC_{i,t}$. Greenness regressors ($\tilde{\beta}^g_{i,t}$ and $\tilde{s}_{i,t}$) are demeaned each month and scaled by fixed (time-invariant) cross-sectional standard deviations as in Section 5. All specifications include the full set of firm controls used in Table 1 (market beta, size, value, and past-return controls), which are omitted for brevity. Coefficients are reported in annual basis points ($10{,}000\times ICC$). Observations are weighted by $1/n_t$ so that each month receives equal total weight; standard errors are clustered by month. Months are weighted equally by observation counts; the pooled coefficient remains an information-weighted average across months when cross-sectional variation differs across months. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels.

A.5 Transitory ESG revisions and green-beta pricing

Our baseline green beta is designed to measure a persistent exposure of a firm's ESG profile to aggregate market greenness, analogous to the role of market beta in the CAPM. Because ESG scores update discretely, a natural concern is that short-horizon rating revisions could induce transitory fluctuations in estimated exposures. If the baseline pricing relation were driven primarily by such high-frequency revisions, the interpretation of green beta as the GSML exposure object would be weaker.

To address this concern, we augment the baseline ICC cross-sections by including both the level green beta and an additional regressor intended to isolate short-horizon ESG revisions (a high-pass proxy). We consider two constructions. First, we estimate a first-difference green beta using monthly changes in firm and market ESG scores. Second, we estimate an AR(1)-residual green beta using AR(1)-filtered ESG innovations. Table A.13 reports joint Fama–MacBeth regressions that include the level green beta together with each transitory proxy.

The level green-beta slope remains negative and statistically significant across specifications. In contrast, the transitory proxies do not retain significance once we condition on market beta, ESG levels, and the full set of firm controls. This pattern indicates that the priced component is the persistent co-movement captured by the baseline green beta rather than short-lived ESG revisions. The sample is slightly shorter than in the main tables because the transitory constructions require consecutive monthly ESG observations.

Table A.13: Joint Fama–MacBeth Regressions of ICC: Level Green Beta and Short-Horizon ESG Revision Proxies

	Spec. A	Spec. B	Spec. C	Spec. D
Panel A: Level green beta and first-difference green beta				
Intercept	0.0853***	0.0664***	0.0824***	0.0671***
	(69.90)	(46.05)	(74.16)	(31.95)
Green beta (level)	-0.0017**	-0.0022**	-0.0016**	-0.0012**
	(-2.54)	(-2.51)	(-1.99)	(-2.07)
First-difference green beta	0.0027*	0.0015	0.0010	0.0004
	(1.87)	(1.27)	(0.89)	(0.36)
Market beta		0.0167***	0.0162***	0.0134***
		(18.72)	(18.82)	(19.27)
ESG score level			-0.0035***	-0.0025***
			(-8.66)	(-8.51)
log(ME)				0.0025***
				(9.14)
log(B/M)				0.0065***
				(16.17)
Short-term reversal (1–0)				-0.0223***
				(-8.46)
Momentum (12–1)				-0.0131***
				(-7.52)
Long-term reversal (60–12)				-0.0050***
				(-6.58)
R^2 (avg.)	0.0041	0.1046	0.1152	0.2328
Months	148	148	148	148
Firm-month observations	136,540	136,540	136,540	136,540
Panel B: Level green beta and AR(1)-innovation green beta				
Intercept	0.0853***	0.0664***	0.0823***	0.0671***
	(69.92)	(46.01)	(73.95)	(31.84)
Green beta (level)	-0.0018***	-0.0023**	-0.0017**	-0.0014**
	(-2.62)	(-2.56)	(-2.08)	(-2.21)
AR(1)-innovation green beta	0.0030**	0.0017	0.0013	0.0006
	(2.05)	(1.54)	(1.20)	(0.65)
Market beta		0.0167***	0.0162***	0.0134***
		(18.72)	(18.82)	(19.26)
ESG score level			-0.0035***	-0.0025***
			(-8.67)	(-8.51)
log(ME)				0.0025***
				(9.08)
log(B/M)				0.0065***
				(16.17)
Short-term reversal (1–0)				-0.0223***
				(-8.45)
Momentum (12–1)				-0.0131***
				(-7.51)
Long-term reversal (60–12)				-0.0050***
				(-6.56)
R^2 (avg.)	0.0040	0.1045	0.1151	0.2328
Months	148	148	148	148
Firm-month observations	136,540	136,540	136,540	136,540

Notes: This table reports time-series averages of monthly Fama–MacBeth cross-sectional slopes where the dependent variable is the implied cost of capital, $ICC_{i,t}$, measured in annual decimal units. Each panel jointly includes the baseline (level) green beta and a proxy for transitory green-beta variation. Panel A uses a first-difference green beta; Panel B uses an AR(1)-innovation green beta. Specifications follow the nested design in the main ICC tables: Spec. B adds market beta; Spec. C adds the ESG score level; Spec. D adds the standard firm controls (log size, log book-to-market, and past-return controls). Newey–West t-statistics with six lags are in parentheses. Coefficients can be multiplied by 10,000 to obtain annual basis points. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.6 Leave-one-out market ESG benchmark

A potential concern in constructing green beta is mechanical self-inclusion, since the value-weighted market ESG benchmark is formed from the same cross-section of firm scores. To address this concern, we recompute the benchmark using a leave-one-out procedure that excludes firm i from the market aggregate in each month, and we re-estimate green betas accordingly. Appendix Table A.14 shows that the estimated price of green-beta exposure is essentially unchanged across specifications, including when conditioning on ESG score levels and the full control set. This evidence indicates that the baseline estimates are not driven by firm self-inclusion in the benchmark used to construct exposure.

Table A.14: Leave-One-Out Market ESG Benchmark: Fama–MacBeth Regressions of ICC

	Dependent variable: $ICC_{i,t}$			
	Spec A	Spec B	Spec C	Spec D
Intercept	0.0855***	0.0668***	0.0816***	0.0663***
	(71.77)	(47.00)	(79.91)	(32.74)
Green beta (LOO benchmark)	-0.0012***	-0.0018***	-0.0013**	-0.0012**
	(-2.75)	(-2.65)	(-2.17)	(-2.41)
Market beta		0.0163***	0.0159***	0.0132***
		(18.36)	(18.46)	(19.13)
ESG score level			-0.0032***	-0.0023***
			(-8.77)	(-8.88)
log(ME)				0.0025***
				(9.41)
log(B/M)				0.0065***
				(16.32)
Short-term reversal (1–0)				-0.0221***
				(-8.62)
Momentum (12–1)				-0.0130***
				(-7.69)
Long-term reversal (60–12)				-0.0051***
				(-7.28)
R^2 (avg.)	0.0020	0.1005	0.1101	0.2278
Months	152	152	152	152
Firm-month observations	144,275	133,969	133,969	126,520

Notes: This table repeats the baseline ICC Fama–MacBeth specifications, but constructs the market ESG benchmark using a leave-one-out procedure that excludes firm i from the value-weighted market ESG aggregate in month t. Green beta is estimated using the same rolling-window procedure as in the baseline and then standardized using the fixed-scale convention. Spec B adds market beta; Spec C adds the ESG score level; Spec D adds the full set of firm controls. Reported coefficients are time-series averages of monthly cross-sectional slopes. Newey–West t-statistics are in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

A.7 Alternative Greenness Proxy: Environmental Pillar Score

This appendix replaces the aggregate ESG score with MSCI's Environmental Pillar Score as an alternative greenness proxy. We (i) recompute firms' green betas using the environmental-score factor and (ii) replace the ESG level regressor with the environmental pillar score level in the cross-sectional ICC regressions. This design isolates whether the priced component documented in the main analysis reflects a dynamic exposure channel (green beta) rather than the level of the greenness characteristic.

Table A.15 reports stock-level Fama–MacBeth regressions of ICC using the environmental pillar score. The environmental green beta remains negatively priced and statistically significant across specifications, including when conditioning on the environmental score level and the full set of controls, whereas the environmental score level provides limited incremental explanatory power once controls are included. Tables A.16 and A.17 examine state-dependent pricing with climate-concern shocks (MCCC), showing that the time variation is concentrated in the price of environmental green beta rather than the environmental score channel. Finally, Tables A.18–A.19 summarize time dependence using a normalized trend and a pre-/post-2016 split.

Table A.15: Fama–MacBeth Regressions of ICC (Environmental Pillar Score)

	Dependent variable: $ICC_{i,t}$			
	Spec A	Spec B	Spec C	Spec D
Intercept	0.0857***	0.0672***	0.0685***	0.0577***
	(72.02)	*(46.22)*	*(53.55)*	*(26.00)*
Environmental green beta	−0.0006***	−0.0006***	−0.0005***	−0.0005***
	(-4.91)	*(-4.35)*	*(-3.51)*	*(-4.14)*
Market beta		0.0162***	0.0160***	0.0132***
		(18.14)	*(18.18)*	*(18.60)*
Environmental pillar score (level)			−0.0003***	0.00005
			(-2.97)	*(0.48)*
log(ME)				0.0023***
				(7.77)
log(B/M)				0.0068***
				(16.39)
Momentum (12–1)				−0.0218***
				(-8.48)
Long-term reversal (60–12)				−0.0050***
				(-7.14)
$\widehat{R}^2$	0.003	0.099	0.101	0.223
Months	152	152	152	152
Firm-month observations	144,286	133,976	133,976	126,527

Notes. The dependent variable is $ICC_{i,t}$ in annual decimal units. Environmental green beta is constructed analogously to the baseline green beta, but using MSCI's Environmental Pillar Score as the greenness proxy in Eq. (5.2). The environmental score level is the contemporaneous pillar score. Greenness regressors are demeaned each month and scaled by fixed cross-sectional standard deviations (time-series averages of within-month dispersions), so coefficients can be interpreted as changes in annual ICC per one-standard-deviation difference on a constant scale. Each monthly cross-section is estimated by OLS with equal firm weights. Newey–West t-statistics with six lags are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.16: State Dependence of Monthly Prices of Environmental Green Beta and Environmental Score (MCCC Shocks)

	Panel A: Price of environmental green beta, λ_t^g (bps per unit β^g)			
	(1) Linear	(2) Quadratic	(3) Sign split	(4) High/Low
Intercept	-5.28***	-4.81***	-4.34***	-5.16***
	(-5.66)	(-4.69)	(-3.63)	(-4.79)
Shock_{t-1}	1.60**	1.81**		
	(2.36)	(2.54)		
Shock_{t-1}^2		-0.47		
		(-1.19)		
Shock_{t-1}^+			0.59	
			(0.60)	
Shock_{t-1}^-			2.95**	
			(1.99)	
High shock_{t-1}				1.97
				(1.47)
Low shock_{t-1}				-3.13
				(-1.20)
	Panel B: Price of environmental score, π_t^s (bps per unit score)			
Intercept	0.46	-0.30	-0.63	0.10
	(0.64)	(-0.37)	(-0.63)	(0.13)
Shock_{t-1}	-0.22	-0.55		
	(-0.36)	(-0.85)		
Shock_{t-1}^2		0.75*		
		(1.79)		
Shock_{t-1}^+			0.94	
			(1.04)	
Shock_{t-1}^-			-1.77	
			(-1.29)	
High shock_{t-1}				0.10
				(0.06)
Low shock_{t-1}				2.36
				(1.23)
Months	152	152	152	152

Notes. This table reports time-series regressions of monthly price series estimated from ICC cross-sections in which the greenness proxy is MSCI's Environmental Pillar Score. $Shock_{t-1}$ is the standardized MCCC climate-concern shock lagged by one month. Coefficients are in annual basis points (ICC units multiplied by 10,000). Newey–West t-statistics are in parentheses. In the sign-split specification, $Shock^+ = Shock \cdot \mathbb{1}\{Shock > 0\}$ and $Shock^- = Shock \cdot \mathbb{1}\{Shock \leq 0\}$. In the high/low specification, high/low shocks are regime indicators constructed from the standardized shock series (e.g., top/bottom tercile).

Table A.17: State Dependence of the Implied Q5–Q1 Environmental Green-Beta Premium

	(1) $Prem_t^{Q5-Q1}$ (bps)
Intercept	-20.29***
	(-8.24)
Shock_t	0.82
	(0.44)
Shock_{t-1}	4.04*
	(1.94)
Months	152

Notes. The dependent variable is the implied environmental Q5–Q1 premium in annual basis points, constructed as $Prem_t^{Q5-Q1} = 10{,}000 \times \widehat{\lambda}_t^g \times (\beta_{Q5,t}^g - \beta_{Q1,t}^g)$ from the environmental-pillar specification. Shocks are standardized MCCC series entered contemporaneously and lagged. Newey–West t-statistics are in parentheses.

Table A.18: Time Trend in Environmental Green-Beta Pricing Objects

	(1) g_t^β (bps per 1σ β^g)	(2) g_t^s (bps per 1σ score)	(3) $Prem_t^{Q5-Q1}$ (bps)	(4) $\Delta\beta_t^g$ (units)
g_1 (Intercept)	-10.86***	-1.76	-30.15***	4.82***
	(-5.49)	(-0.64)	(-5.61)	(6.34)
g_2 ($\text{Trend}_{[0,1]}$)	6.73**	7.18	19.58**	2.98***
	(2.05)	(1.56)	(2.22)	(3.01)
Months	152	152	152	152

Notes. This table reports time-trend regressions of monthly pricing objects estimated from the environmental-pillar ICC cross-sections. The trend variable is normalized to the unit interval. g_t^β and g_t^s are annual basis-point prices of environmental green beta and the environmental score level, respectively. $Prem_t^{Q5-Q1}$ is the implied annual Q5–Q1 premium in basis points, and $\Delta\beta_t^g$ is the within-month Q5–Q1 spread in environmental green beta (in beta units). Newey–West t-statistics are in parentheses.

Table A.19: Pre-2016 vs Post-2016 Comparison (Environmental Pillar Score)

	(1) λ_t^g (bps per unit β^g)	(2) $Prem_t^{Q5-Q1}$ (bps)	(3) $\Delta\beta_t^g$ (units)	(4) π_t^s (bps per unit score)
Pre-2016 mean (Intercept)	-9.01***	-26.26***	5.82***	-0.72
	(-4.78)	(-6.74)	(8.13)	(-0.62)
Post-2016 shift (Post2016)	6.70***	10.68**	0.89	2.14
	(3.46)	(2.14)	(1.23)	(1.48)
Months (total)	152	152	152	152

Notes. This table reports a pre-/post-2016 split for environmental-pillar pricing objects. The "Pre-2016 intercept" is the mean prior to January 2016 in the estimation sample; the "Post-2016 shift" is the difference between the post-2016 mean and the pre-2016 mean. Prices and premia are in annual basis points; $\Delta\beta_t^g$ is in beta units. Newey–West t-statistics are in parentheses.

Table A.20: Pooled Month Fixed-Effects Regressions (Equal-Weighted Months): Environmental Pillar Score

	(1) Linear	(2) Quadratic	(3) Sign split	(4) High/Low
$\beta^g_{i,t}$	-6.99***	-6.89***	-7.02***	-7.22***
	(-9.06)	(-7.73)	(-5.87)	(-8.01)
Score level	0.73	1.23	2.19	1.84
	(0.71)	(0.97)	(1.34)	(1.54)
$\beta^g_{i,t} \times Shock_{t-1}$	1.75**	1.79**		
	(2.22)	(2.26)		
$\beta^g_{i,t} \times Shock^2_{t-1}$		-0.08		
		(-0.16)		
Score$\times Shock_{t-1}$	-1.82*	-1.60		
	(-1.87)	(-1.55)		
Score$\times Shock^2_{t-1}$		-0.49		
		(-0.71)		
$\beta^g_{i,t} \times Shock^+_{t-1}$			1.84	
			(1.29)	
$\beta^g_{i,t} \times Shock^-_{t-1}$			1.63	
			(1.04)	
Score$\times Shock^+_{t-1}$			-3.38*	
			(-1.98)	
Score$\times Shock^-_{t-1}$			0.28	
			(0.14)	
$\beta^g_{i,t} \times HighShock_{t-1}$				3.00
				(1.41)
$\beta^g_{i,t} \times LowShock_{t-1}$				-1.54
				(-0.64)
Score$\times HighShock_{t-1}$				-7.16**
				(-2.59)
Score$\times LowShock_{t-1}$				0.13
				(0.05)
Observations	126,527	126,527	126,527	126,527
Months (FE)	152	152	152	152
Within R^2	0.185	0.185	0.185	0.185

Notes. This table reports pooled regressions with month fixed effects using environmental-pillar greenness measures. Environmental green beta and the environmental score level are standardized following the paper's fixed-scale convention. Columns correspond to alternative nonlinear/state-dependence specifications interacted with the lagged standardized MCCC shock. Coefficients are in annual basis points (ICC units multiplied by 10,000). Month fixed effects are included and months are weighted equally; standard errors are clustered by month.

A.8 Alternative Greenness Proxy: Trucost Scope 1 + 2 Emissions Intensity

This appendix replaces the aggregate ESG score with a Trucost-based greenness proxy constructed from firms' Scope 1 + 2 emissions intensity. We (i) recompute firms' green betas using the Trucost-score factor and (ii) replace the ESG level regressor with the Trucost greenness score level in the cross-sectional ICC regressions. This design isolates whether the priced component documented in the main analysis reflects a dynamic exposure channel (green beta) rather than the level of the greenness characteristic.

Table A.21 reports stock-level Fama–MacBeth regressions of ICC using the Trucost proxy. The Trucost greenness score is oriented so that higher values correspond to *lower* Scope 1 + 2 intensity (higher = greener). Unlike MSCI scores which update at (roughly) monthly frequency, Trucost intensity measures are observed annually and carried forward to monthly observations, which may introduce additional measurement noise and attenuate estimated relations. In the Fama–MacBeth regressions, the Trucost green beta enters with a negative but statistically weak coefficient, same for the Trucost greenness level once included.

Tables A.22 and A.23 examine state-dependent pricing with climate-concern shocks (MCCC). Finally, Tables A.24–A.25 summarize time dependence using a normalized trend and a pre-/post-2016 split. Table A.26 provides a pooled month-FE robustness analogue.

Table A.21: Fama–MacBeth Regressions of ICC (Trucost Scope 1 + 2 Intensity Greenness)

	Dependent variable: $ICC_{i,t}$			
	Spec A	Spec B	Spec C	Spec D
Intercept	0.0858***	0.0672***	0.0668***	0.0604***
	(72.91)	*(39.48)*	*(38.73)*	*(19.46)*
Trucost green beta	-0.0091*	-0.0078	-0.0039	-0.0039
	(-1.77)	*(-1.49)*	*(-1.13)*	*(-1.57)*
Market beta		0.0166***	0.0166***	0.0125***
		(15.11)	*(14.95)*	*(13.55)*
Trucost greenness score (level)			0.0002	0.0019*
			(0.15)	*(1.78)*
log(ME)				0.0021***
				(6.96)
log(B/M)				0.0070***
				(16.81)
Short-term reversal (1–0)				-0.0225***
				(-7.79)
Momentum (12–1)				-0.0132***
				(-7.47)
Long-term reversal (60–12)				-0.0050***
				(-6.58)
$\widehat{R}^2$	0.0067	0.1110	0.1151	0.2420
Months	150	150	150	150
Firm-month observations	111,440	104,060	104,060	99,507

Notes. The dependent variable is $ICC_{i,t}$ in annual decimal units. Trucost green beta is constructed analogously to the baseline green beta, but using Trucost Scope 1 + 2 intensity as the greenness measure (firm and market series) in the green-beta construction. The Trucost score level is the contemporaneous Scope 1 + 2 intensity level. Greenness regressors are demeaned each month and scaled by fixed cross-sectional standard deviations (time-series averages of within-month dispersions), so coefficients can be interpreted as changes in annual ICC per one-standard-deviation difference on a constant scale. Each monthly cross-section is estimated by OLS with equal firm weights. Newey–West t-statistics are in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.22: State Dependence of Monthly Prices of Trucost Green Beta and Trucost Greenness (MCCC Shocks)

	Panel A: Price of Trucost green beta, λ_t^g (bps per unit β^g)			
	(1) Linear	(2) Quadratic	(3) Sign split	(4) High/Low
Intercept	-11.88	-13.12	-11.62	-4.09
	(-1.62)	(-1.45)	(-0.86)	(-0.57)
Shock_{t-1}	8.91	8.44		
	(0.88)	(0.70)		
Shock^2_{t-1}		1.26		
		(0.18)		
Shock^+_{t-1}			8.61	
			(0.91)	
Shock^-_{t-1}			9.27	
			(0.33)	
High shock_{t-1}				0.84
				(0.12)
Low shock_{t-1}				-56.06
				(-0.95)
	Panel B: Price of Trucost greenness score, π_t^s (bps per unit score)			
Intercept	2.62*	0.06	-0.50	2.26**
	(1.85)	(0.04)	(-0.23)	(2.25)
Shock_{t-1}	-1.86	-2.82		
	(-0.80)	(-0.94)		
Shock^2_{t-1}		2.62		
		(1.10)		
Shock^+_{t-1}			1.62	
			(1.01)	
Shock^-_{t-1}			-6.29	
			(-0.91)	
High shock_{t-1}				-0.66
				(-0.56)
Low shock_{t-1}				3.56
				(0.61)
Months	150	150	150	150

Notes. This table reports time-series regressions of monthly price series estimated from ICC cross-sections that use Trucost-based greenness measures. $Shock_{t-1}$ is the standardized MCCC climate-concern shock lagged by one month. Coefficients are in annual basis points (ICC units multiplied by 10,000). Newey–West t-statistics are in parentheses. In the sign-split specification, $Shock^+ = Shock \cdot \mathbb{1}\{Shock > 0\}$ and $Shock^- = Shock \cdot \mathbb{1}\{Shock \leq 0\}$. In the high/low specification, high/low shocks are regime indicators constructed from the standardized shock series (e.g., top/bottom tercile).

Table A.23: State Dependence of the Implied Q5–Q1 Trucost Green-Beta Premium

	(1) $Prem_t^{Q5-Q1}$ (bps)
Intercept	-9.31***
	(-2.59)
Shock_t	-1.62
	(-0.62)
Shock_{t-1}	-1.02
	(-0.38)
Months	150

Notes. The dependent variable is the implied Trucost Q5–Q1 premium in annual basis points, constructed as $Prem_t^{Q5-Q1} = 10{,}000 \times \widehat{\lambda}_t^g \times (\beta_{Q5,t}^g - \beta_{Q1,t}^g)$ from the Trucost specification. Shocks are standardized MCCC series entered contemporaneously and lagged. Newey–West t-statistics are in parentheses.

Table A.24: Time Trend in Trucost Green-Beta Pricing Objects

	(1) g_t^β (bps per 1σ β^g)	(2) g_t^s (bps per 1σ score)	(3) $Prem_t^{Q5-Q1}$ (bps)	(4) $\Delta\beta_t^g$ (units)
g_1 (Intercept)	-6.36***	17.84***	-13.61***	3.62***
	(-3.27)	(4.03)	(-3.07)	(8.94)
g_2 ($\text{Trend}_{[0,1]}$)	1.45	11.35	2.45	1.62**
	(0.35)	(1.38)	(0.25)	(2.61)
Months	150	150	150	150

Notes. This table reports time-trend regressions of monthly pricing objects estimated from Trucost-based ICC cross-sections. The trend variable is normalized to the unit interval. g_t^β and g_t^s are annual basis-point prices of Trucost green beta and the Trucost score level, respectively. $Prem_t^{Q5-Q1}$ is the implied annual Q5–Q1 premium in basis points, and $\Delta\beta_t^g$ is the within-month Q5–Q1 spread in Trucost green beta (in beta units). Newey–West t-statistics are in parentheses.

Table A.25: Pre-2016 vs Post-2016 Comparison (Trucost Scope 1+2 Intensity Greenness)

	(1) λ_t^g (bps per unit β^g)	(2) $Prem_t^{Q5-Q1}$ (bps)	(3) $\Delta\beta_t^g$ (units)	(4) π_t^s (bps per unit score)
Pre-2016 mean (Intercept)	-1.46	-10.81***	3.82***	13.12***
	(-0.85)	(-4.58)	(14.50)	(5.47)
Post-2016 shift (Post2016)	-1.41	-2.81	1.10***	1.50
	(-0.77)	(-0.71)	(3.82)	(0.49)
Months (total)	150	150	150	150

Notes. This table reports a pre-/post-2016 split for Trucost-based pricing objects. The "Pre-2016 intercept" is the mean prior to January 2016 in the estimation sample; the "Post-2016 shift" is the difference between the post-2016 mean and the pre-2016 mean. Prices and premia are in annual basis points; $\Delta\beta_t^g$ is in beta units. Newey–West t-statistics are in parentheses.

Table A.26: Pooled Month Fixed-Effects Regressions (Equal-Weighted Months): Trucost Scope 1+2 Intensity Greenness

	(1) Linear	(2) Quadratic	(3) Sign split	(4) High/Low
$\beta^g_{i,t}$	-5.85***	-4.48***	-2.90*	-3.41***
	(-4.85)	(-3.36)	(-1.69)	(-2.63)
Score level (greenness)	14.06***	14.21***	14.86***	15.09***
	(14.02)	(13.93)	(10.61)	(14.57)
$\beta^g_{i,t} \times Shock_{t-1}$	-2.00	-1.41		
	(-1.33)	(-0.96)		
$\beta^g_{i,t} \times Shock^2_{t-1}$		-2.35		
		(-1.57)		
Score$\times Shock_{t-1}$	0.06	0.24		
	(0.07)	(0.29)		
Score$\times Shock^2_{t-1}$		-0.15		
		(-0.18)		
$\beta^g_{i,t} \times Shock^+_{t-1}$			-2.18	
			(-1.36)	
$\beta^g_{i,t} \times Shock^-_{t-1}$			-2.12	
			(-0.73)	
Score$\times Shock^+_{t-1}$			-0.04	
			(-0.03)	
Score$\times Shock^-_{t-1}$			1.48	
			(0.44)	
$\beta^g_{i,t} \times HighShock_{t-1}$				-0.47
				(-0.32)
$\beta^g_{i,t} \times LowShock_{t-1}$				-9.80
				(-1.23)
Score$\times HighShock_{t-1}$				0.40
				(0.46)
Score$\times LowShock_{t-1}$				1.58
				(0.55)
Observations	99,534	99,534	99,534	99,534
Months (FE)	152	152	152	152
Within R^2	0.194	0.194	0.194	0.194

Notes. This table reports pooled regressions with month fixed effects using Trucost-based greenness measures. Trucost green beta and the Trucost score level are standardized following the paper's fixed-scale convention. Columns correspond to alternative nonlinear/state-dependence specifications interacted with the lagged standardized MCCC shock. Coefficients are in annual basis points (ICC units multiplied by 10,000). Month fixed effects are included and months are weighted equally; standard errors are clustered by month.

Table A.27: Variable definitions and constructions

Variable	Definition and construction
$\mathrm{ICC}_{i,t}$	Implied cost of capital for firm i in month t (annual, decimal units), computed from analysts' forecasts and current price using standard ICC implementations (see text for model details).
$\mathrm{ESG}_{i,t}$	MSCI ESG score for firm i in month t (scaled as in the paper; higher values indicate stronger ESG profile).
$s_{m,t}$	Market-wide ESG level in month t, computed as the value-weighted mean ESG score across firms, using a leave-one-out scheme when constructing firm-level exposures (see text).
$\beta^{g}_{i,t}$ (Green beta)	Estimated from a rolling regression of firm i's ESG score on the market-wide ESG series: $\mathrm{ESG}_{i,\tau} = \alpha_{i,t} + \beta^{g}_{i,t} s_{m,\tau} + \varepsilon_{i,\tau}$ over a rolling window (e.g., 36 months), with standard filtering and winsorization as described in the paper.
$\mathrm{ME}_{i,t}$	Market equity (market capitalization). In the characteristics data, ME is constructed as price times shares outstanding (reported in million USD).
$\mathrm{B/M}_{i,t}$	Book-to-market equity. In the characteristics data, B/M is constructed as book equity divided by market equity.
$\beta^{m}_{i,t}$ (Market beta)	60-month CAPM beta estimated as $\mathrm{Cov}_{60}(\mathrm{RET}_{i,t}, \mathrm{MKT_RF}_t)/\mathrm{Var}_{60}(\mathrm{MKT_RF}_t)$.
$\mathrm{ret}_{1,0}$	Short-term reversal (1-month return): $\mathrm{RI}_t/\mathrm{RI}_{t-1} - 1$.
$\mathrm{ret}_{12,1}$	Momentum (12–1): $\mathrm{RI}_{t-1}/\mathrm{RI}_{t-12} - 1$.
$\mathrm{ret}_{60,12}$	Long-run reversal (60–12): $\mathrm{RI}_{t-12}/\mathrm{RI}_{t-60} - 1$.
MCCC_t	Media Climate Change Concern (MCCC) index (monthly).
Shock_t	Climate-salience shock defined as the monthly innovation in log(MCCC) and standardized to mean zero and unit variance (see text).

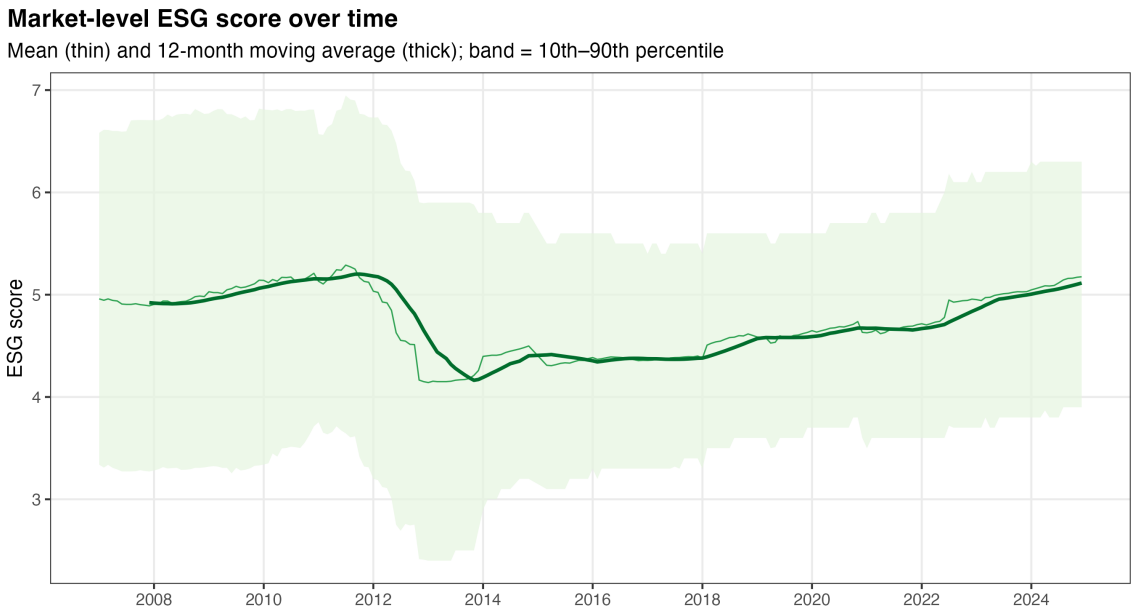


Figure A.1: Market-level ESG score over time. The thin line plots the monthly cross-sectional mean raw ESG score across firms and the thick line is the 12-month moving average. The shaded band shows the 10th–90th percentile range across firms in each month.

A.9 Figures

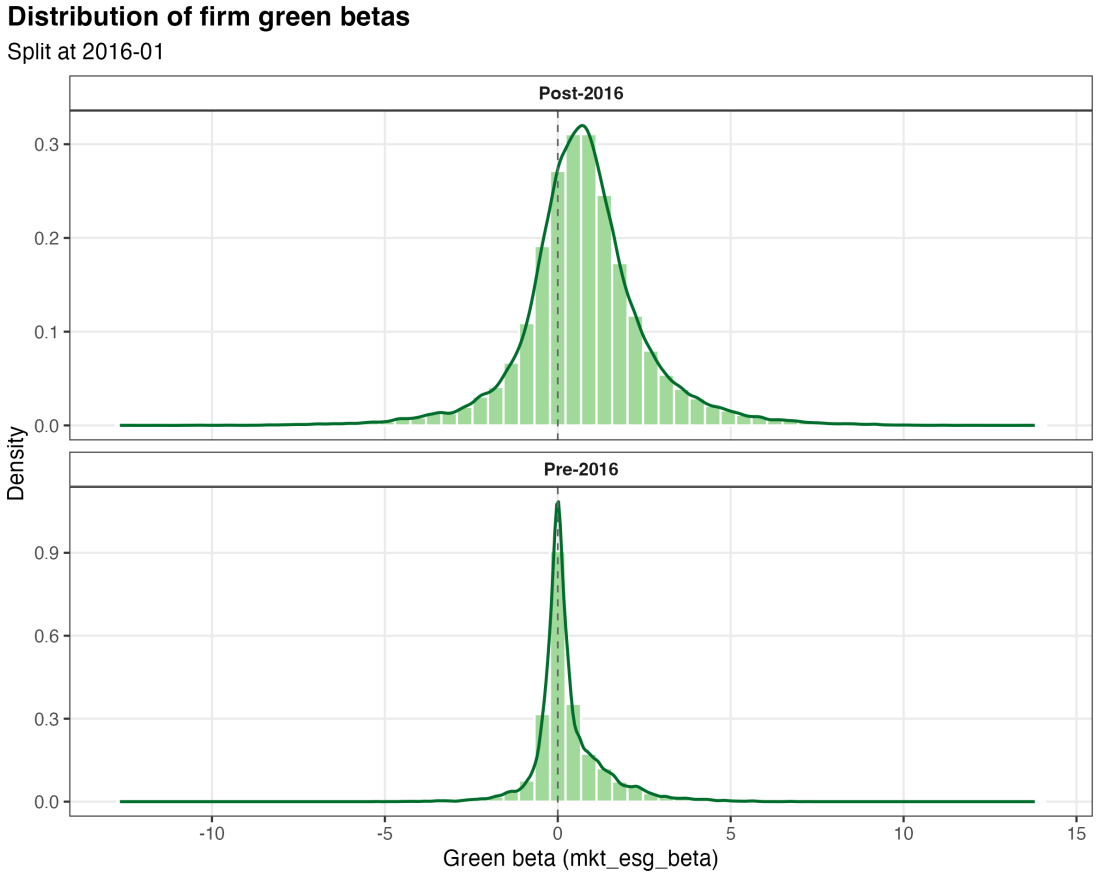


Figure A.2: Distribution of firm green betas. The figure shows histograms and kernel densities of the firm-level green beta $\hat{\beta}^{g}_{i,t}$, split into pre-2016 and post-2016 subsamples. The vertical dashed line marks zero.

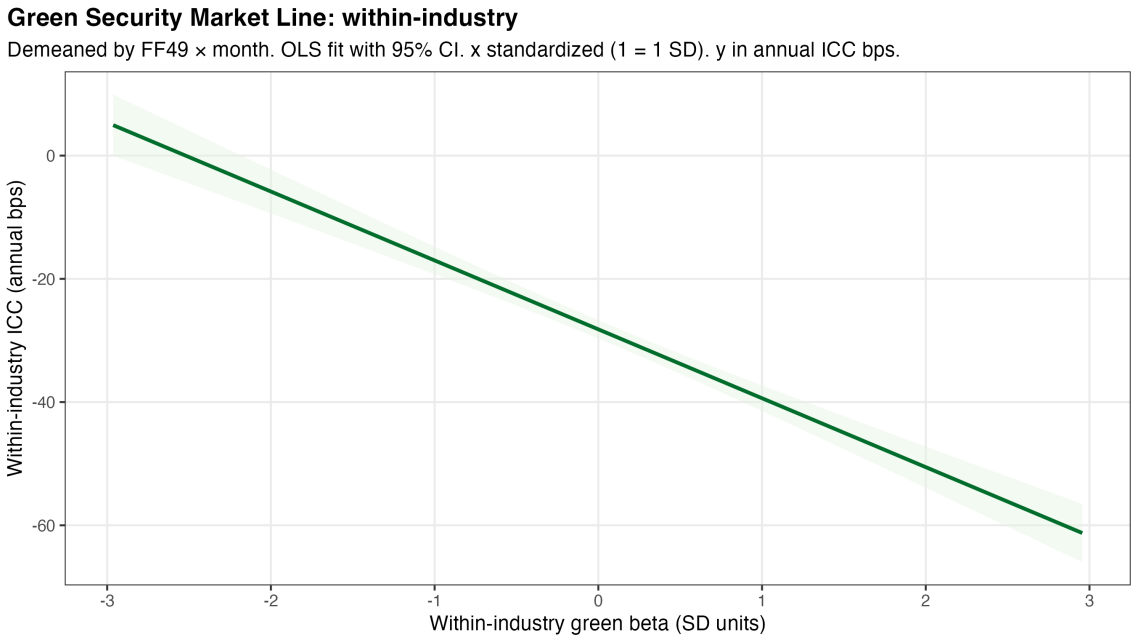


Figure A.3: Green Security Market Line: within-industry relationship. The figure plots the fitted OLS relation between within-industry ICC and within-industry green beta after demeaning both variables by FF49 industry × month. Green beta is standardised (1 unit = 1 cross-sectional standard deviation). The y-axis is in annual ICC basis points; the shaded band is a 95% confidence interval.

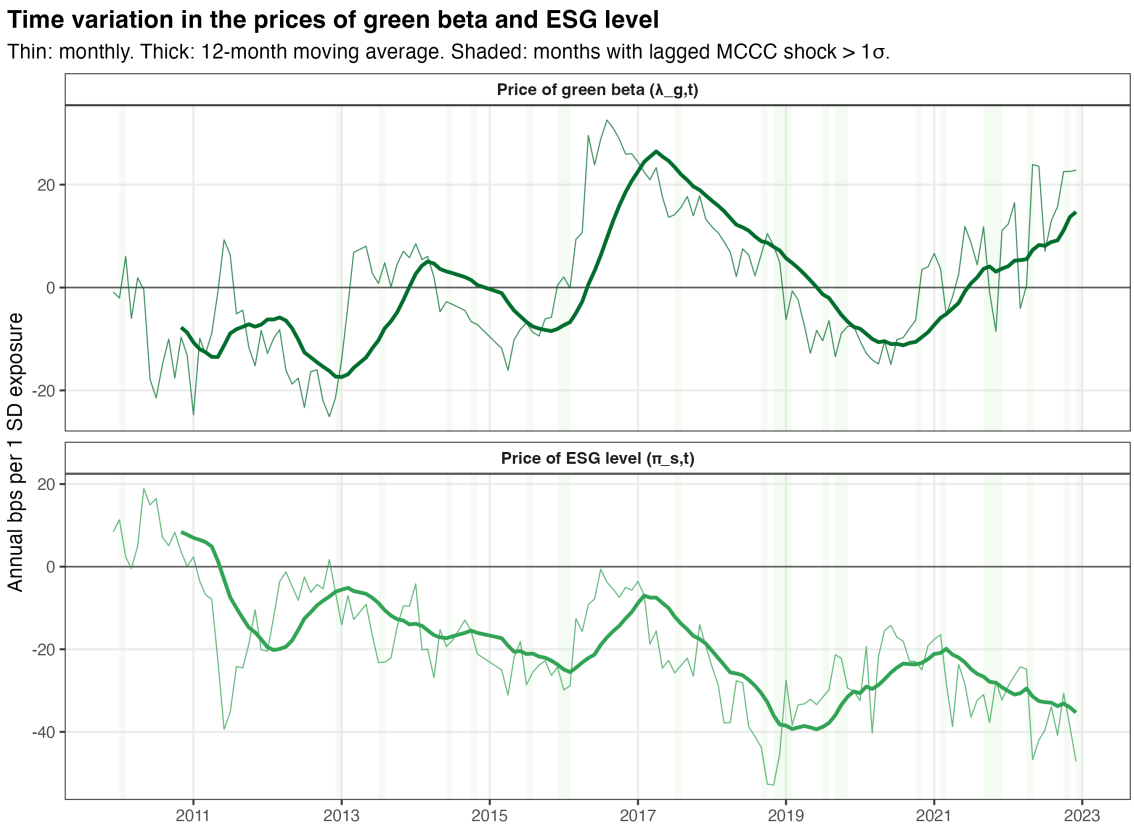


Figure A.4: Time variation in the prices of green-beta exposure and ESG-score level (monthly series). Thin lines plot the monthly estimates of λ_t^{β} and π_t^{s}; thick lines are 12-month moving averages. Series are expressed in annual basis points per one-standard-deviation difference in the relevant characteristic. Shaded regions denote climate-concern months, defined as months in which the lagged standardised MCCC innovation exceeds 1σ.

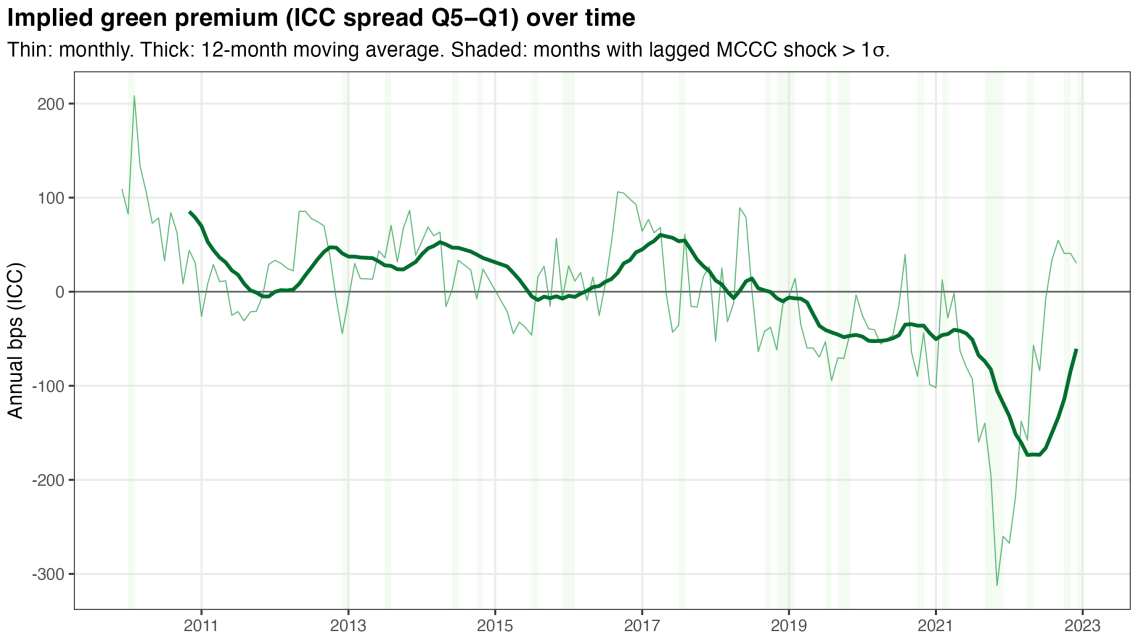


Figure A.5: Implied green premium over time. The figure plots the time series of the ICC spread between green-beta quintile 5 and quintile 1 portfolios (Q5–Q1), expressed in annual basis points. The thin line is the monthly series and the thick line is the 12-month moving average. Shaded regions denote climate-concern months (lagged standardised MCCC innovation $> 1\sigma$).

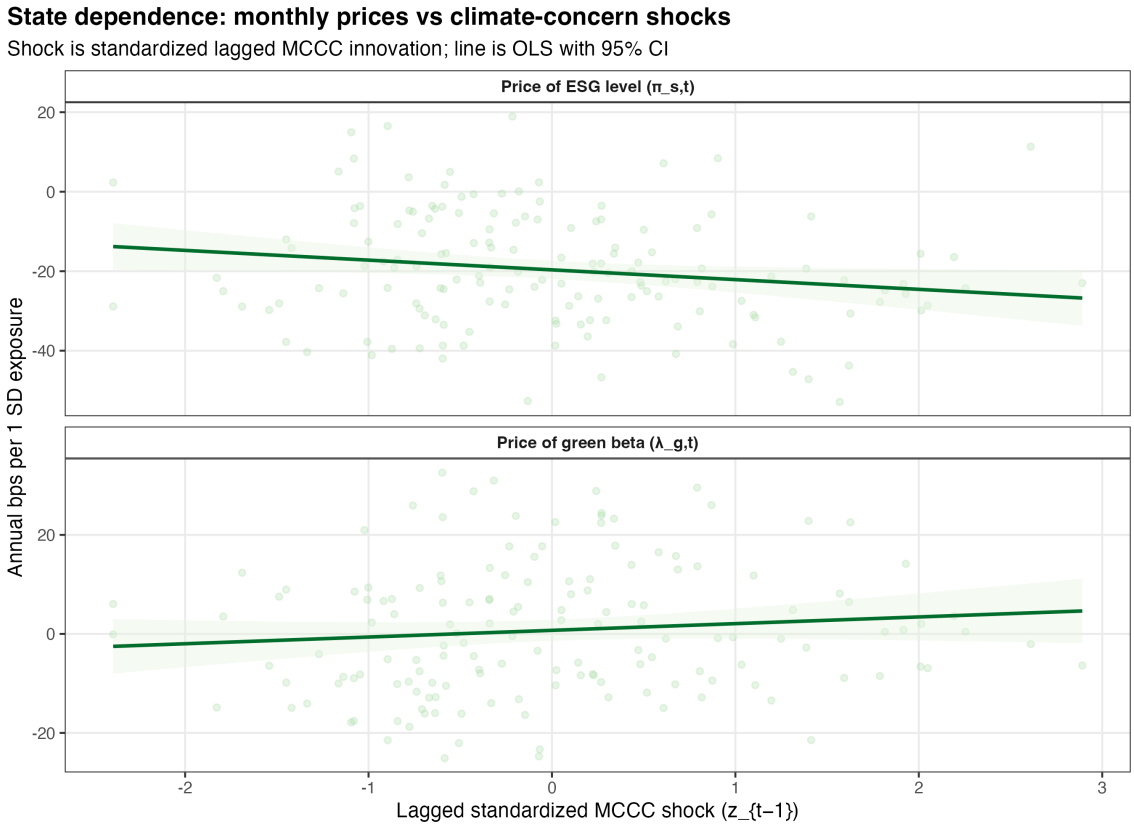


Figure A.6: State dependence: monthly prices vs climate-concern shocks (scatter). Each point is a month. The x-axis is the lagged standardised MCCC innovation and the y-axis is the monthly price (annual basis points per one-standard-deviation exposure). The solid line is the fitted OLS relation with a 95% confidence band.

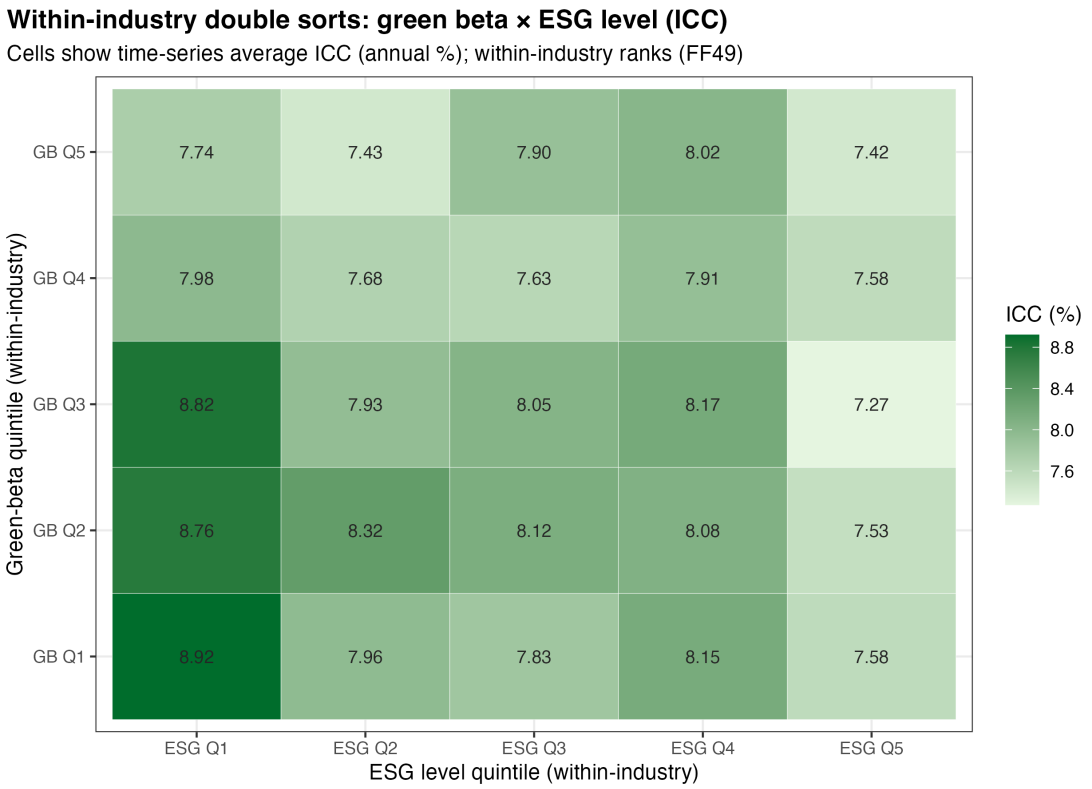


Figure A.7: Within-industry double sorts: green beta × ESG level (ICC). The figure reports time-series average ICC (annual %) for 5×5 portfolios formed using within-industry ranks (FF49) of lagged ESG level and lagged green beta. Cell values are portfolio means.