



ISSN: 2288-7709

JEMM website: <https://accesson.kr/jemm>doi: <http://doi.org/10.20482/jemm.2026.14.1.69>

Is Carbon Risk Priced? Evidence from Korean Stock Markets

Dong-Jin PYO¹

Received: January 14, 2026. Revised: January 31, 2026. Accepted: February 06, 2026.

Abstract

Purpose: This paper examines whether corporate carbon intensity is priced in the cross-section of stock returns using firm-level data from the Korean equity market. The analysis is motivated by growing investor attention to climate transition risks and by mixed empirical evidence on carbon risk pricing in the existing literature. **Research design and methodology:** We develop a Bayesian hierarchical asset-pricing framework that allows the pricing of carbon intensity to vary across industries while controlling for firm-specific characteristics, unobserved firm-level heterogeneity, and aggregate time effects. The hierarchical structure enables partial pooling across sectors and firms, while a heavy-tailed return specification accommodates the empirical distribution of stock returns. **Results:** The posterior distribution of the average carbon-intensity coefficient is positive but economically modest, with substantial uncertainty and a credible interval that includes zero, indicating limited evidence of a strong uniform carbon risk premium across firms. At the same time, posterior means of industry-specific carbon effects are predominantly positive, suggesting cross-sectoral variation in how carbon exposure is priced, albeit with considerable uncertainty at the individual industry level. **Conclusions:** Taken together, the results imply that carbon risk is not priced uniformly in the Korean stock market. Instead, carbon exposure is associated with a weak and heterogeneous return premium that operates through sector-dependent channels rather than as a pervasive market-wide effect. These findings underscore the importance of accounting for heterogeneity and uncertainty when assessing the financial implications of firms' carbon intensity.

Keywords: Carbon risk, Asset pricing, Climate finance, Bayesian hierarchical models

JEL Classification Code: G12, Q54, C11, C23, G11

1. Introduction

Climate change has emerged as a central source of economic and financial risk, affecting firms through physical damages, regulatory interventions, technological transitions, and shifts in consumer demand. As governments strengthen climate-related policies and investors increasingly incorporate environmental considerations into valuation and portfolio decisions, a fundamental question in asset pricing is whether firms' exposure to carbon-related risks is reflected in expected stock returns. Understanding how carbon emissions are priced in financial markets is particularly relevant for economies with carbon-intensive industrial structures and evolving regulatory frameworks,

such as South Korea. Despite its recent commitments to carbon neutrality, South Korea remains heavily reliant on energy-intensive manufacturing and export-driven growth, creating a setting in which climate transition risks are salient but their pricing in financial markets remains uncertain.

A growing body of research examines whether climate-related risks constitute a priced factor in equity markets. Early contributions emphasize that firms differ substantially in their exposure to climate transition risk through regulatory uncertainty, stranded-asset risk, and adjustment costs associated with decarbonization, thereby providing a theoretical basis for potential compensation in expected returns (Hong et al., 2020). More recent studies provide direct empirical evidence linking carbon emissions to stock

¹ First Author. Associate Professor, Department of Economics, Changwon National University, Republic of Korea, Email: djpyo@changwon.ac.kr

© Copyright: The Author(s)

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>) which permits unrestricted noncommercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

returns. Using U.S. data, Bolton and Kacperczyk (2021) document that firms with higher carbon emissions earn higher expected returns, interpreting this pattern as compensation for carbon risk rather than mispricing. Related work further shows that carbon-intensive firms face higher costs of capital and exhibit heightened sensitivity to climate policy shocks and regulatory developments (Pankratz et al., 2023).

Despite this growing literature, empirical findings remain mixed, particularly outside the U.S. context. Several studies report weak, unstable, or economically small average carbon risk premia, suggesting that the pricing of carbon exposure may depend on institutional environments, market structures, and industry composition (Ilhan et al., 2023; Pankratz et al., 2023). International evidence further indicates that the strength and direction of carbon risk pricing are shaped by regulatory credibility and cross-sectoral differences in production technologies and energy dependence (Hong et al., 2020; Hong et al., 2019). These findings raise the possibility that carbon risk is priced heterogeneously across industries rather than through a uniform market-wide premium.

Two empirical challenges complicate the identification of carbon risk premia. First, carbon exposure varies systematically across industries due to differences in production processes, energy inputs, and regulatory exposure. Treating carbon emissions as a uniformly priced firm characteristic may therefore obscure meaningful cross-sectoral heterogeneity. Second, stock returns exhibit substantial idiosyncratic volatility and heavy tails, which can attenuate economically relevant but modest pricing effects when Gaussian return models are imposed. Together, these features motivate econometric approaches that explicitly account for heterogeneity and distributional non-normality.

This paper addresses these challenges by investigating how firm-level carbon emissions are priced in the cross-section of Korean stock returns using a Bayesian hierarchical panel framework. Rather than imposing a common carbon premium across all firms, we allow the marginal effect of carbon intensity on expected returns to vary across industries through sector-specific random slopes. Firm-level random intercepts capture persistent unobserved heterogeneity in average returns, while year effects absorb aggregate market-wide shocks. To accommodate the heavy-tailed nature of stock returns, return innovations are modeled using a Student- t likelihood, which improves robustness and predictive performance relative to Gaussian specifications.

Our empirical findings can be summarized as follows. First, the average carbon risk premium is positive but economically modest and imprecisely estimated. The posterior distribution of the aggregate carbon-intensity

coefficient is centered slightly above zero, while its credible interval includes zero, indicating limited evidence of a strong uniform carbon premium across firms. Second, we document sectoral heterogeneity in carbon risk pricing. Although individual sector-specific carbon premia are estimated with substantial uncertainty, posterior means are predominantly positive across industries, and the estimated dispersion of industry-level effects is non-negligible. This pattern suggests that carbon exposure tends to be priced in a similar direction across sectors, albeit weakly and heterogeneously. Third, complementary portfolio simulations reveal that firms with higher carbon intensity earn higher average returns than low-emission firms, particularly over longer holding horizons, consistent with carbon exposure being associated with risk compensation rather than transitory mispricing.

This study contributes to the literature in several ways. Methodologically, it demonstrates the value of hierarchical Bayesian models for analyzing climate-related risks in asset pricing environments characterized by noisy data and complex heterogeneity. Empirically, it provides new evidence on carbon risk pricing in South Korea, an economy that has received relatively limited attention in the climate finance literature despite its carbon-intensive industrial structure. More broadly, the results highlight that carbon risk may be priced weakly but pervasively, operating through sector-dependent channels rather than as a uniform market-wide premium.

The remainder of the paper is organized as follows. Section 2 presents the empirical model and estimation strategy. Section 3 describes the data and variable construction. Section 4 reports the main regression results and sectoral heterogeneity analysis. Section 5 examines portfolio implications through carbon-sorted trading strategies. Section 6 concludes.

2. Methodology

2.1. Empirical Specification

To examine whether carbon-related transition risk is priced in the cross-section of stock returns, we estimate a hierarchical Bayesian panel model that explicitly allows for heterogeneity in the pricing of carbon exposure across industries, while controlling for firm-level characteristics and aggregate time effects.

Let $R_{i,t}$ denote the stock return of firm i in year t . The empirical model specification is given by

$$R_{i,t} = \beta_0 + (\beta_c + \alpha_{s(i)}) CI_{i,t} + \boldsymbol{\beta}'\mathbf{X}_{i,t} + \delta_i + \lambda_t + \varepsilon_{i,t} \quad (1)$$

where $CI_{i,t}$ denotes firm i 's carbon intensity, defined as carbon emissions scaled by sales, $\mathbf{X}_{i,t}$ is a vector of firm-level control variables, δ_i captures unobserved firm-specific heterogeneity, λ_t represents year effects, and $\varepsilon_{i,t}$ is an idiosyncratic error term.

The key feature of equation 1 is that the effect of carbon intensity on returns is allowed to vary across industries. Specifically, β_c represents the average pricing of carbon intensity across all firms, while $\alpha_{s(i)}$ captures industry-specific deviations from the average effect, where $s(i)$ indexes the industry to which firm i belongs. This structure allows the data to determine whether carbon risk is priced heterogeneously across sectors, reflecting differences in technology, regulation, and exposure to the low-carbon transition.

2.2. Hierarchical Structure

To account for unobserved heterogeneity across sectors, firms, and time, we employ a hierarchical modeling structure with partial pooling. This multilevel specification decomposes variation in stock returns into sector-level, firm-level, and aggregate time-series components, improving identification and mitigating overfitting in an unbalanced panel.

Sector-level heterogeneity in the pricing of carbon intensity is modeled through sector-specific random slopes. The sector-specific deviation from the average carbon–return relationship is given by

$$\alpha_s = \sigma_\alpha \tilde{\alpha}_s, \quad \tilde{\alpha}_s \sim \mathcal{N}(0,1), \quad s = 1, \dots, S \quad (2)$$

where σ_α governs the dispersion of heterogeneity in carbon risk pricing across sectors. This structure allows the carbon–return relationship to vary systematically across sectors while shrinking estimates toward the common effect when sector-specific information is limited.

Firm-level unobserved heterogeneity is captured by firm-specific random intercepts,

$$\delta_i = \sigma_\delta \tilde{\delta}_i, \quad \tilde{\delta}_i \sim \mathcal{N}(0,1), \quad i = 1, \dots, N \quad (3)$$

which account for persistent differences in average returns across firms due to time-invariant characteristics. Partial pooling across firms is particularly important given heterogeneous observation counts and reduces noise-driven estimates for firms with short return histories.

Aggregate time-series variation in expected returns is accommodated through year-specific random effects,

$$\lambda_t \sim \mathcal{N}(0, \sigma_{\text{year}}^2), \quad t = 1, \dots, T \quad (4)$$

which absorb common shocks affecting all firms in a given year, such as macroeconomic conditions or market-wide risk sentiment.

All hierarchical components are implemented using non-centered parameterizations to enhance computational efficiency and improve convergence. Together, this structure enables a coherent decomposition of return variation across sectors, firms, and time, while allowing for flexible heterogeneity in the pricing of carbon exposure.

2.3. Priors and Distributional Assumptions

We adopt weakly informative priors that regularize estimation while allowing the data to play a dominant role in identifying both average effects and heterogeneous components. All regressors are standardized prior to estimation to improve numerical stability and facilitate comparison of coefficient magnitudes across covariates.

Let $\boldsymbol{\beta}$ denote the vector of slope coefficients associated with the standardized regressors, and let β_0 denote the intercept. We assign diffuse Normal priors,

$$\begin{aligned} \beta_0 &\sim \mathcal{N}(0, 10^2), \\ \boldsymbol{\beta} &\sim \mathcal{N}(\mathbf{0}, 10^2 \mathbf{I}) \end{aligned} \quad (5)$$

reflecting prior uncertainty regarding the magnitude and direction of covariate effects. Given the standardized scale of the regressors, these priors are weakly informative and impose minimal restrictions on economically plausible parameter values.

As the baseline specification, idiosyncratic stock return innovations are modeled using a Student- t distribution to accommodate fat tails commonly observed in asset returns:

$$\varepsilon_{i,t} \sim t_\nu(0, \sigma) \quad (6)$$

The scale parameter is assigned a Half-Normal prior,

$$\sigma \sim \text{Half-Normal}(2 s_\gamma) \quad (7)$$

where s_γ denotes the sample standard deviation of annual stock returns, anchoring the prior to the empirical scale of the dependent variable. The degrees-of-freedom parameter ν is estimated from the data subject to a lower bound ensuring finite variance, allowing the likelihood to flexibly adapt to tail thickness.

The standard deviations of the hierarchical components are assigned Half-Normal priors,

$$\begin{aligned} \sigma_\alpha &\sim \text{Half-Normal}(0.5), \\ \sigma_\delta &\sim \text{Half-Normal}(0.5), \\ \sigma_{\text{year}} &\sim \text{Half-Normal}(1), \end{aligned} \quad (8)$$

which provide moderate shrinkage and stabilize estimation in the presence of many groups and unbalanced data.

Overall, the combination of a heavy-tailed likelihood, weakly informative priors, and hierarchical pooling yields a flexible yet well-regularized modeling framework that is well suited to the empirical characteristics of firm-level stock return data.

2.4. Control Variables

To isolate the pricing effect of carbon intensity from other well-established determinants of stock returns, we include a comprehensive set of firm-level and market-level control variables. All firm-specific controls are measured at $t - 1$ to mitigate potential simultaneity and look-ahead bias, while market-wide variables are measured contemporaneously. This specification follows standard practice in empirical asset-pricing studies using panel data.

Firm size is controlled for using the logarithm of total assets, $\log(\text{Assets}_{i,t-1})$. Firm size is one of the most robust predictors of expected returns and captures systematic differences in information environment, liquidity, and risk exposure across firms. Smaller firms have historically earned higher average returns, a phenomenon commonly referred to as the size effect (Banz, 1981; Fama & French, 1992). Using total assets rather than market capitalization is standard in firm-level panel settings that include multiple balance-sheet controls.

In the model, financial leverage is captured by the ratio of total liabilities to total assets, $\text{Liabilities}/\text{Assets}_{i,t-1}$. Leverage is positively related to expected returns due to higher financial distress risk and equity volatility (Bhandari, 1988). Importantly, leverage is also correlated with carbon intensity, as capital-intensive and carbon-intensive firms often rely more heavily on debt financing. Controlling for leverage therefore helps disentangle carbon risk from balance-sheet-driven risk premia.

Corporate liquidity and financial flexibility are controlled for using the cash-to-assets ratio, $\text{Cash}/\text{Assets}_{i,t-1}$. Firms with higher cash holdings tend to face lower financing constraints and are better positioned to absorb regulatory or transition-related shocks. Prior studies document that cash holdings are a key determinant of firm value, risk-taking behavior, and resilience to adverse shocks (Bates et al., 2009; Opler et al., 1999). Including liquidity controls is particularly relevant in a climate-risk context, where adjustment costs may be substantial.

Profitability is measured using return on assets (ROA), defined as net income divided by total assets. Profitability is a first-order determinant of expected returns and is closely linked to firms' operating efficiency and investment opportunities. Empirical asset-pricing models consistently

find that more profitable firms earn higher average returns (Fama & French, 2015; Novy-Marx, 2013). Controlling for profitability prevents return premia associated with operating performance from being confounded with carbon exposure.

Firm growth and investment activity are proxied by lagged sales growth, $\text{SalesGrowth}_{i,t-1}$. Firms experiencing rapid expansion often exhibit different risk profiles and expected returns compared to mature firms. The investment-based asset-pricing literature shows that firms with aggressive growth or investment tend to earn lower subsequent returns (Cooper et al., 2008; Fama & French, 2015). Because expanding firms may simultaneously increase production and emissions, controlling for growth is especially important when studying carbon intensity.

To account for aggregate market movements, we include the contemporaneous annual return on the KOSPI index. Market returns capture systematic shocks that affect all firms simultaneously and ensure that estimated carbon effects are not driven by broad market cycles. This control complements the inclusion of year effects and improves identification of firm-specific pricing effects.

We also control for short-run return persistence by including lagged firm returns, $R_{i,t-1}$. Empirical evidence shows that past returns can predict future returns due to investor underreaction and momentum effects (Jegadeesh & Titman, 1993). Including lagged returns helps separate carbon-related pricing from return dynamics unrelated to fundamentals.

Firm valuation characteristics are captured using lagged price-to-book (PBR $_{i,t-1}$) and price-to-earnings (PER $_{i,t-1}$) ratios. These variables proxy for value versus growth orientation and expected cash-flow prospects. The value premium, commonly measured using book-to-market or its inverse, is a central feature of asset-pricing models (Fama & French, 1992, 1993). Including both valuation ratios allows for flexible control of valuation effects that may correlate with carbon intensity.

Finally, we include each firm's market beta with respect to the KOSPI index estimated from historical return data. Beta captures exposure to systematic market risk in the Capital Asset Pricing Model (CAPM) framework (Sharpe, 1964). Controlling for beta ensures that the estimated carbon premium is not merely compensation for higher covariance with aggregate market returns.

Overall, this set of controls reflects standard practice in empirical asset-pricing studies and allows us to attribute variation in expected returns more cleanly to carbon intensity and its heterogeneous pricing across industries.

2.5. Posterior Computation

The model is estimated using Bayesian Markov Chain Monte Carlo (MCMC) methods. Posterior inference is based on draws from the joint posterior distribution obtained via the No-U-Turn Sampler (NUTS), an adaptive variant of Hamiltonian Monte Carlo that is well suited for high-dimensional hierarchical models (Hoffman & Gelman, 2014). All models are estimated using multiple chains with dispersed initial values. Convergence is assessed using standard diagnostics, including the potential scale reduction factor (\hat{R}) and effective sample sizes, following established best practices in Bayesian computation (Gelman et al., 2013; Vehtari et al., 2021).

The Bayesian hierarchical framework employed in this study offers several advantages relative to conventional fixed-effects panel regressions. First, partial pooling across industries and firms improves estimation efficiency and mitigates overfitting in the presence of unbalanced panels and heterogeneous observation counts (Gelman & Hill, 2007). Second, the hierarchical structure enables direct probabilistic inference on cross-sectoral heterogeneity in carbon risk pricing, rather than relying on ad hoc interaction terms or sequential estimation approaches. Third, the Bayesian framework naturally propagates parameter uncertainty into posterior summaries, allowing for coherent inference on both average effects and sector-specific deviations (Rossi et al., 2005).

These features are particularly valuable in the context of climate-related asset pricing, where true effects are likely to be economically modest, heterogeneous, and estimated with substantial uncertainty. By explicitly modeling heterogeneity and uncertainty, the Bayesian approach provides a flexible and transparent framework for assessing the pricing of carbon exposure in financial markets.

3. Data

3.1. Sample Construction

The empirical analysis is conducted using an annual panel dataset of Korean listed firms. The initial sample consists of all non-financial firms listed on the Korea Exchange (KRX) for which financial statement information, stock return data, and carbon emissions data are available. Financial firms are excluded due to their distinct balance-sheet structures and regulatory environments.¹

Firm-level financial statement data are obtained from the NICE Value Search, while stock return and market information are collected from the Korea Exchange. Firm-level carbon emissions data are drawn from the National

Greenhouse Gas Management System (NGMS), which provides verified annual emissions for entities subject to Korea's mandatory greenhouse gas reporting requirements. Because the original NGMS records do not contain firm-level identifiers, emissions linked to business registration numbers are obtained through Eco & Partners, an ESG consulting firm specializing in environmental data integration. Owing to limitations in firm-identifier availability, the matched emissions data cover the period from 2012 to 2020.² Observations with missing values in the dependent variable or key regressors are excluded. All variables are winsorized at the 1st and 99th percentiles to reduce the impact of extreme observations.

3.2. Carbon Intensity Measure and Controls

The primary variable of interest is firm-level carbon intensity. Carbon intensity is defined as the ratio of total carbon emissions to total sales:³

$$\text{Carbon Intensity}_{i,t} = \frac{\text{Emissions}_{i,t}}{\text{Sales}_{i,t}} \quad (9)$$

To address potential simultaneity between emissions and stock returns and to reflect the information set available to investors, carbon intensity is lagged by one year in the specification. This lag structure is standard in the climate-finance literature and ensures that carbon exposure is predetermined with respect to subsequent stock returns.

Following the asset-pricing literature, we include a comprehensive set of firm-level control variables that are known to affect expected stock returns. All firm characteristics are measured at year $t - 1$. Firm size is measured as the natural logarithm of total assets. Financial leverage is proxied by the ratio of total liabilities to total assets, capturing financial distress risk. Liquidity and financial flexibility are measured using the cash-to-assets ratio. Profitability is captured by return on assets (ROA), defined as net income divided by total assets. Firm growth is proxied by sales growth.

To account for return dynamics and valuation effects, we include lagged stock returns, the lagged price-to-book ratio (PBR), and the lagged price-to-earnings ratio (PER). These variables control for momentum, value-versus-growth characteristics, and expected cash-flow prospects. Finally, firm-level market beta is included to control for systematic risk exposure.

All explanatory variables are standardized prior to estimation to improve numerical stability in the Bayesian hierarchical model and to facilitate comparison of coefficient magnitudes. Stock returns are expressed in annual terms and are standardized where appropriate for estimation purposes.

Descriptive statistics for the main variables are reported in Table 1. The table indicates substantial cross-sectional heterogeneity in firms' carbon intensity and financial characteristics. Carbon intensity exhibits considerable dispersion, with a mean of 0.324 and a long right tail, suggesting that a small subset of firms accounts for disproportionately high emissions relative to sales. Firm size, measured by log assets, also varies widely across the sample,

reflecting the presence of both small and large firms. Financial characteristics such as leverage, liquidity, profitability, and growth display meaningful variation, while stock return measures and valuation ratios exhibit substantial volatility. Overall, the pronounced dispersion across both carbon exposure and firm fundamentals provides a suitable empirical setting for examining heterogeneous pricing of carbon-related risks across firms and industries.

Table 1: Summary Statistics

Variable	Mean	Std.	Min	25%	Median	75%	Max
Carbon intensity	0.324	0.607	0.004	0.053	0.138	0.368	7.194
Log assets	21.111	1.771	18.133	19.722	20.623	22.313	26.486
Leverage (Liabilities/Assets)	0.496	0.190	0.091	0.350	0.552	0.635	0.935
Cash ratio (Cash/Assets)	0.060	0.045	0.001	0.025	0.050	0.084	0.251
Return on assets (ROA)	2.391	4.857	-19.750	0.290	2.350	5.020	18.720
Sales growth	4.405	16.642	-42.140	-4.955	1.970	10.493	109.120
KOSPI return	0.061	0.132	-0.173	0.007	0.036	0.094	0.308
Lagged stock return	0.159	0.946	-0.952	-0.186	-0.016	0.231	12.080
Price-to-book ratio	1.067	0.899	0.141	0.524	0.803	1.257	6.651
Price-to-earnings ratio	13.589	36.640	-232.502	4.243	12.182	22.407	207.418
Market beta	0.853	0.378	-0.011	0.568	0.851	1.127	1.766

Notes: This table reports summary statistics for the variables used in the empirical analysis. Carbon intensity is defined as carbon emissions scaled by sales and is lagged by one year. All firm-level control variables are measured at $t - 1$, while market variables are measured contemporaneously. Financial variables are based on IFRS financial statements.

4. Results

This section presents the empirical results from the Bayesian hierarchical model described in Section 2. We report results from the specification that allows for industry-level heterogeneity in the pricing of carbon intensity, while controlling for firm characteristics, aggregate market conditions, and year effects. Posterior inference is based on the full joint distribution of model parameters, enabling a probabilistic assessment of both average pricing effects and cross-sectional heterogeneity.

4.1. Overall Carbon Risk Premium

Table 2 reports posterior summary statistics for the baseline model, including posterior means, standard deviations, and highest-density intervals (HDIs) for the main regression coefficients and variance components. All parameters exhibit satisfactory convergence, with \hat{R} statistics close to unity, indicating stable posterior sampling.

The posterior mean of the average carbon-intensity coefficient is positive but small, and its 95% HDI includes zero. This indicates that, after controlling for firm

characteristics, market conditions, and time effects, there is no strong evidence of a precisely estimated uniform carbon risk premium across firms. The concentration of posterior mass near zero suggests that any average pricing effect of carbon intensity is economically modest and subject to substantial uncertainty.

Figure 1 displays the posterior distribution of the average carbon-intensity coefficient, β_{carbon} . The posterior mean is approximately 0.01, implying a slightly positive association between carbon intensity and expected stock returns on average. However, the posterior distribution is tightly centered around zero, and the 95% HDI spans both negative and positive values. This pattern indicates that, while the data provide directional support for a positive average relationship between carbon exposure and returns, the magnitude of this effect is small and not precisely identified. Overall, the posterior evidence does not support a strong or pervasive average carbon risk premium operating uniformly across firms.

A natural interpretation for the weak and imprecisely estimated average carbon premium relates to the institutional design and implementation of Korea's Emissions Trading Scheme (K-ETS) during the sample

period. In contrast to more mature carbon markets, such as the EU Emissions Trading System (EU ETS), the K-ETS was characterized for much of the period by extensive free allocation of allowances, relatively modest compliance stringency, and limited price discovery in secondary permit markets. In particular, a large share of emissions permits under the K-ETS was allocated without charge, especially in the early and middle phases of the program. As a result, many firms faced low effective marginal costs of emissions, at least up to their allocated allowance levels. Even when firms exceeded free allocations, permit prices remained comparatively low and exhibited limited volatility relative to those observed in the EU ETS. These institutional features are well documented in comparative policy assessments, which show that Korea's effective carbon prices—after accounting for free allocation and coverage—remained substantially below those observed in the EU ETS during most of the sample period (Organization for Economic Co-Operation and Development, 2023).

These institutional features imply that carbon pricing risk may have been weakly binding for a broad cross-section of firms, reducing the extent to which carbon intensity translated into meaningful expected cash-flow or cost-of-capital impacts at the aggregate level. In such an environment, carbon intensity is less likely to represent a pervasive, market-wide risk factor that commands a strong average return premium. Instead, the pricing of carbon exposure may operate primarily through more localized channels—such as firms operating near compliance thresholds, firms with limited abatement flexibility, or firms concentrated in sectors subject to tighter regulatory scrutiny.

More broadly, these results suggest that the pricing of carbon risk depends not only on firms' emissions profiles but also on the credibility, stringency, and market depth of the underlying regulatory framework. When carbon pricing mechanisms are relatively mild or predictable, as in Korea during much of the sample period, carbon exposure may be priced weakly and heterogeneously rather than as a strong uniform risk premium. This institutional context helps reconcile the absence of a precisely estimated average carbon premium with the presence of economically meaningful cross-sectional and long-horizon return differentials. Importantly, this interpretation does not imply that carbon risk is irrelevant for asset prices. Rather, it highlights that the financial salience of carbon exposure is endogenous to policy design and enforcement. As carbon markets mature, allowance allocations tighten, and permit prices become more volatile, the average pricing of carbon intensity may strengthen over time—a possibility that future research could explore as Korea's climate policy regime evolves.

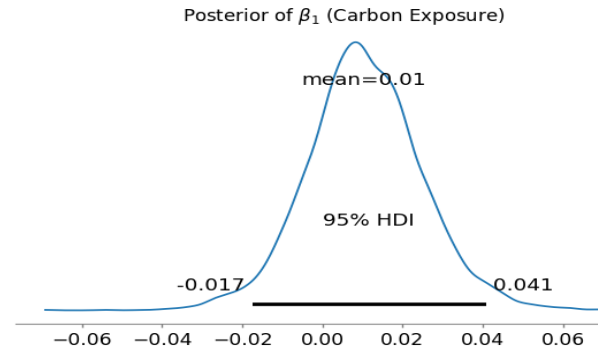


Figure 1: Posterior Distribution of the Average Carbon-Intensity Coefficient

Notes: The figure shows the posterior density of the average carbon-intensity coefficient β_{carbon} from the baseline Bayesian hierarchical model. The horizontal line indicates the 95% highest-density interval (HDI), and the vertical marker denotes the posterior mean.

While the average carbon-intensity coefficient is weakly identified, the estimated dispersion of industry-specific deviations, summarized by σ_{α} , is positive but modest. As reported in Table 2, the posterior mean of σ_{α} is approximately 0.019, with its posterior mass concentrated close to zero. This finding suggests limited, though non-negligible, heterogeneity in the pricing of carbon intensity across industries. At the same time, individual industry-level deviations α_s are estimated with substantial uncertainty, and their posterior intervals generally overlap zero. Taken together, these results indicate that carbon risk is not priced uniformly across sectors, but that cross-industry differences in the carbon–return relationship are economically small and imprecisely estimated at the individual industry level.

Firm-level unobserved heterogeneity, captured by the random intercept variance σ_{δ} , is estimated precisely and is substantially larger than the corresponding industry-level variance component. The posterior mean of σ_{δ} is approximately 0.08, with a tight posterior interval, indicating pronounced and persistent differences in average returns across firms that are not explained by observable characteristics or industry affiliation. The residual scale parameter σ is also tightly estimated, reflecting the substantial idiosyncratic volatility characteristic of annual stock returns, even after accounting for firm-, industry-, and time-level effects.

Several control variables exhibit posterior distributions that are meaningfully different from zero. Firm size, measured by log assets, is negatively associated with expected returns, consistent with the size effect documented in the asset pricing literature. Leverage and sales growth display positive posterior means with HDIs that exclude zero, indicating compensation for financial risk and growth-

related uncertainty. Cash holdings are also positively associated with returns, though the estimated magnitude is modest. Market-wide conditions, proxied by contemporaneous KOSPI returns, load positively on firm-level returns, confirming the importance of aggregate market movements even after controlling for year effects.

In contrast, profitability, valuation ratios (PBR and PER), lagged returns, and market beta exhibit posterior distributions centered close to zero, with HDIs that include zero. This suggests limited incremental explanatory power of these variables once firm characteristics, market conditions, and hierarchical effects are taken into account.

Overall, the baseline results indicate that carbon intensity does not command a strong or precisely estimated average return premium. While there is evidence of modest cross-industry heterogeneity in the pricing of carbon exposure, the magnitude of these differences is small and subject to considerable uncertainty. The findings therefore point toward weak and heterogeneous channels through which carbon exposure may affect expected returns, rather than a pervasive or uniform carbon risk premium across firms.

Table 2: Posterior Summary Statistics for the Model

Parameters	mean	sd	hdi_3%	hdi_97%	\hat{R}
beta[carbon intensity]	0.0104	0.0145	-0.0152	0.0400	1.0006
beta[log assets]	-0.0237	0.0101	-0.0423	-0.0047	1.0002
beta[leverage]	0.0398	0.0107	0.0196	0.0589	1.0005
beta[cash ratio]	0.0185	0.0091	0.0013	0.0354	1.0004
beta[ROA]	-0.0079	0.0100	-0.0270	0.0103	1.0003
beta[sales growth]	0.0469	0.0120	0.0247	0.0690	1.0006
beta[KOSPI]	0.0674	0.0243	0.0226	0.1125	1.0007
beta[lagged stock return]	0.0068	0.0078	-0.0079	0.0215	1.0002
beta[PBR]	-0.0026	0.0054	-0.0130	0.0075	1.0003
beta[PER]	0.0004	0.0042	-0.0077	0.0084	1.0001
beta[beta_KOSPI]	0.0112	0.0079	-0.0038	0.0261	1.0003
σ	0.2134	0.0026	0.2085	0.2184	1.0002
σ_α	0.0189	0.0127	0.0006	0.0417	1.0007
σ_δ	0.0821	0.0046	0.0736	0.0907	1.0008

Notes: This table reports posterior means, standard deviations, and 94% highest-density intervals (HDI) from the Bayesian hierarchical model with a Student-*t* likelihood. All regressors are standardized. Sector-specific coefficients α_s capture deviations from the average carbon–return slope (reported separately). Estimation is based on four MCMC chains with convergence assessed using the \hat{R} statistic.

4.2. Sectoral Carbon Risk Premium

To further examine heterogeneity in the pricing of carbon exposure, we analyze the posterior distributions of sector-specific carbon premia, defined as the sum of the average effect and the sectoral deviation, $\beta_{\text{carbon}} + \alpha_s$. These quantities capture how the marginal effect of carbon intensity on expected returns varies across industries.

The posterior summaries reported in Table 3 reveal modest cross-sector dispersion in estimated carbon premia. Posterior means differ across industries and are predominantly positive, with most estimates lying in a narrow range close to zero. At the same time, the corresponding highest-density intervals are wide and

generally overlap zero, indicating substantial uncertainty in the estimation of individual sectoral effects.

Despite this uncertainty, an important pattern emerges from the posterior distributions. Across the majority of sectors, the posterior means of $\beta_{\text{carbon}} + \alpha_s$ are positive, suggesting that higher carbon intensity is, on average, associated with higher expected returns within sectors. This directional consistency appears even though sector-level effects are not precisely estimated in a statistical sense. Thus, the evidence points toward a weak but broadly positive carbon premium that operates across sectors.

The dispersion of posterior means across sectors is consistent with the estimated industry-level variance component, σ_α , which is positive but modest in magnitude.

This evidence indicates that carbon exposure is not priced identically across industries, while also suggesting that cross-sector differences in carbon pricing are economically small relative to firm-level heterogeneity and overall idiosyncratic return variation. Taken together, the results point to a weak but broadly positive carbon premium that varies across sectors, operating with limited strength and considerable uncertainty rather than as a sharply differentiated industry-specific effect.

Importantly, the lack of precisely estimated sector-specific premia cautions against strong conclusions regarding which industries face systematically higher or lower carbon risk premia. Instead, the results suggest that

carbon-related risks are priced in a similar direction across sectors, with industry-specific factors introducing noise and estimation uncertainty rather than clear-cut segmentation in carbon pricing.

Overall, the sectoral posterior evidence supports a view in which carbon exposure commands a weak and heterogeneous premium that is broadly positive in expectation but imprecisely estimated at the industry level. This interpretation is consistent with the portfolio-based evidence reported below, where return differentials emerge when firms are sorted on carbon intensity, despite substantial cross-sectional volatility at the individual firm and sector levels.

Table 3: Posterior Summaries of Sector-Specific Carbon Premia ($\beta_{\text{carbon}} + \alpha_s$).

Sector	Mean	SD	HDI _{2.5%}	HDI _{97.5%}	\hat{R}
Retail Trade (Excluding Motor Vehicles)	0.019	0.032	-0.043	0.091	1.00
Beverage Manufacturing	-0.006	0.037	-0.092	0.057	1.00
Land Transport and Pipeline Transport	0.016	0.034	-0.051	0.091	1.00
Professional Services	-0.001	0.037	-0.084	0.069	1.00
Motor Vehicles and Trailer Manufacturing	0.022	0.025	-0.023	0.077	1.00
Non-metallic Mineral Products Manufacturing	0.025	0.032	-0.033	0.095	1.00
Electrical Equipment Manufacturing	0.000	0.036	-0.084	0.065	1.00
Electronic Components, Computers, Audio, Video and Communication Equipment Manufacturing	-0.008	0.035	-0.086	0.051	1.00
Basic Metal Manufacturing	-0.001	0.016	-0.034	0.028	1.00
General Construction	-0.003	0.043	-0.099	0.073	1.00
Textile Products Manufacturing (Excluding Apparel)	0.012	0.016	-0.018	0.044	1.00
Pulp, Paper and Paper Products Manufacturing	0.007	0.028	-0.050	0.069	1.00
Food Manufacturing	0.006	0.018	-0.032	0.039	1.00
Chemical Substances and Chemical Products Manufacturing (Excluding Pharmaceuticals)	0.022	0.032	-0.040	0.088	1.00
Rubber and Plastics Products Manufacturing	0.022	0.039	-0.048	0.111	1.00
Wood and Cork Products Manufacturing (Excluding Furniture)	0.007	0.029	-0.054	0.069	1.00
Wholesale Trade and Commodity Brokerage	0.007	0.035	-0.067	0.086	1.00
Air Transport	0.008	0.038	-0.071	0.089	1.00
Warehousing and Transport Support Services	0.020	0.029	-0.031	0.085	1.00
Electricity, Gas, Steam and Air Conditioning Supply	0.006	0.036	-0.075	0.078	1.00
Leather, Bags and Footwear Manufacturing	0.012	0.030	-0.051	0.075	1.00
Financial Services	0.010	0.037	-0.072	0.087	1.00
Fabricated Metal Products Manufacturing (Excluding Machinery and Furniture)	0.014	0.035	-0.055	0.095	1.00
Other Transport Equipment Manufacturing	0.007	0.035	-0.072	0.076	1.00
Coke, Briquettes and Refined Petroleum Products Manufacturing	0.023	0.039	-0.050	0.111	1.00
Postal and Telecommunications Services	-0.008	0.040	-0.104	0.057	1.00

Computer Programming, Systems Integration and Management Services	0.001	0.040	-0.087	0.078	1.00
Other Machinery and Equipment Manufacturing	0.020	0.040	-0.049	0.117	1.00
Medical, Precision and Optical Instruments and Watches Manufacturing	0.016	0.036	-0.052	0.101	1.00
Insurance and Pension Funding	0.037	0.048	-0.036	0.149	1.00
Tobacco Manufacturing	0.010	0.037	-0.063	0.095	1.00
Sports and Recreation Services	0.010	0.040	-0.073	0.095	1.00
Pharmaceutical Products and Medicinal Materials Manufacturing	0.011	0.040	-0.076	0.095	1.00
Motion Picture, Video and Audio Recording Production and Distribution	0.011	0.040	-0.071	0.094	1.00
Waste Collection, Treatment and Materials Recovery	0.011	0.040	-0.072	0.096	1.00

Notes: This table reports posterior means, standard deviations, and 95% highest-density intervals (HDI) for sector-specific carbon premia, defined as the sum of the average carbon intensity coefficient and the sector-level deviation, $\beta_{\text{carbon}} + \alpha_s$. All estimates are obtained from the baseline Bayesian hierarchical model with a Student- t likelihood for return innovations. Sectoral effects are partially pooled through a hierarchical prior, implying shrinkage toward the common effect when sector-level information is limited. All regressors are standardized prior to estimation. Convergence diagnostics based on the \hat{R} statistic indicate satisfactory convergence for all parameters.

5. Carbon-Momentum Trading Simulation

This section evaluates whether carbon exposure contains exploitable return predictability by implementing a carbon-intensity-sorted long-short trading strategy. The simulation is designed to mirror a standard cross-sectional “characteristics-sorted” portfolio approach widely used in empirical asset pricing, while explicitly incorporating proportional transaction costs and alternative portfolio formation/holding horizons. The objective is not to claim implementable arbitrage, but to assess whether carbon intensity is associated with economically meaningful return differentials in a simple, transparent trading rule.

5.1. Portfolio Formation and Holding Protocol

Let $CI_{i,t}$ denote firm i 's carbon intensity (emissions scaled by sales) observed at the end of year t . In the simulation, portfolio formation in year t uses the lagged carbon intensity measure, $CI_{i,t-1}$, ensuring that the sorting signal is predetermined relative to returns earned over the subsequent holding period.

For each formation year $t \geq 2012$, we sort all firms with non-missing $CI_{i,t-1}$ into ascending order. We then form:

a long portfolio consisting of the K firms with the lowest carbon intensity (the “low-emission” group), and

a short portfolio consisting of the K firms with the highest carbon intensity (the “high-emission” group),

where $K \in \{5,10,15,20,25\}$ is varied in a grid search. If fewer than $2K$ firms are available in a given year, the strategy is not formed for that year.

Portfolios are held for H years, where $H \in \{1,2,3,4,5\}$. For each leg (long and short), we compute firm-level gross

holding-period returns by compounding annual simple returns:

$$R_{i,t \rightarrow t+H-1}^{\text{gross}} = \prod_{\tau=t}^{t+H-1} (1 + R_{i,\tau}) - 1 \quad (10)$$

where $R_{i,\tau}$ denotes the annual simple return of firm i in year τ . The long and short leg returns are computed as equal-weighted averages across firms within each leg.

5.2. Transaction Costs and Net Returns

To account for implementation frictions, we incorporate proportional transaction fees at both entry and exit. Let $P_{i,t}$ denote the year-end adjusted stock price for firm i at formation year t , and $P_{i,t+H-1}$ the year-end adjusted price at the end of the holding horizon. For each position, a proportional fee rate ϕ is applied at both purchase and sale (transaction_fee_rate = 0.0005).

The simulation approximates total fee impact on the holding-period return using the ratio of round-trip fees to the initial purchase price:

$$\text{FeeEffect}_{i,t \rightarrow t+H-1} = \phi \cdot \frac{P_{i,t} + P_{i,t+H-1}}{P_{i,t}} \quad (11)$$

and defines the net holding-period return for each firm as

$$R_{i,t \rightarrow t+H-1}^{\text{net}} = R_{i,t \rightarrow t+H-1}^{\text{gross}} - \text{FeeEffect}_{i,t \rightarrow t+H-1} \quad (12)$$

Firms with missing returns or prices (or non-positive purchase prices) are excluded from the corresponding portfolio leg in that formation year. The net long and short

portfolio returns are then computed as the equal-weighted averages of $R_{t,t \rightarrow t+H-1}^{\text{net}}$ across valid firms in each leg.

Finally, the annualized long-short strategy return for formation year t is defined as:

$$R_t^{\text{LS}}(K, H) = \bar{R}_{\text{Long},t}^{\text{net}}(K, H) - \bar{R}_{\text{Short},t}^{\text{net}}(K, H) \quad (13)$$

where $\bar{R}_{\text{Long},t}^{\text{net}}(K, H)$ and $\bar{R}_{\text{Short},t}^{\text{net}}(K, H)$ denote the net holding-period returns for the long and short legs, respectively.

5.3. Performance Metric and Parameter Grid

For each (K, H) pair, we construct a time series of long-short returns $\{R_t^{\text{LS}}(K, H)\}$ across all feasible formation years and summarize performance using the average long-short return:

$$\bar{R}^{\text{LS}}(K, H) = \frac{1}{|\mathcal{T}_{K,H}|} \sum_{t \in \mathcal{T}_{K,H}} R_t^{\text{LS}}(K, H) \quad (14)$$

where $\mathcal{T}_{K,H}$ denotes the set of years for which the strategy can be formed given data availability and the holding horizon.

We evaluate robustness across portfolio concentration and investment horizon by exploring a grid of $K \in \{5, 10, 15, 20, 25\}$ and $H \in \{1, 2, 3, 4, 5\}$. The resulting matrix of $\bar{R}^{\text{LS}}(K, H)$ values is visualized as a heatmap, where rows correspond to holding periods and columns correspond to portfolio sizes. This visualization provides a compact summary of how the average carbon strategy return varies with portfolio construction choices.

5.4. Interpretation

A positive $\bar{R}^{\text{LS}}(K, H)$ implies that, on average, low-emission firms outperform high-emission firms over the holding horizon net of transaction fees, consistent with a carbon-related return differential.⁴ Conversely, a negative value indicates relative underperformance of low-emission firms. Because this simulation abstracts from borrowing constraints, shorting frictions, and liquidity effects beyond proportional fees, the results should be interpreted as descriptive evidence on carbon-sorted return spreads rather than as a fully implementable trading strategy.

5.5. Carbon-Momentum Trading Simulation Results

Figure 2 summarizes the performance of the carbon-intensity-sorted long-short trading strategy across alternative portfolio sizes and holding horizons. Each cell

reports the average net return of a strategy that is long in low-emission firms and short in high-emission firms, after accounting for proportional transaction costs at portfolio entry and exit.

Several patterns emerge from the simulation. First, the average long-short returns are uniformly negative across all combinations of portfolio size and holding period. This indicates that, on average, portfolios of low-emission firms underperform portfolios of high-emission firms over the examined horizons. The magnitude of underperformance varies substantially with the holding period, suggesting that the carbon-related return differential is sensitive to investment horizon.

Second, the negative return differential becomes more pronounced as the holding period lengthens. While one-year holding strategies yield modestly negative average returns, the magnitude of the long-short spread increases substantially for holding periods of three years or longer. For example, strategies with a three-year holding horizon exhibit average returns below -0.30 across most portfolio sizes, with even larger negative values observed for longer horizons. This pattern suggests that any relative underperformance of low-emission firms accumulates over time rather than dissipating quickly.

Third, portfolio concentration plays a secondary but non-negligible role. For a given holding period, increasing the number of firms in each leg tends to moderate the magnitude of the negative return, consistent with diversification reducing the impact of extreme firm-level outcomes. However, the overall sign of the strategy return remains negative regardless of portfolio size, indicating that the result is not driven solely by a small subset of firms.

Taken together, the simulation results imply that a simple carbon-intensity-based trading strategy does not generate positive abnormal returns over the sample period once transaction costs are taken into account. Instead, the negative long-short returns are consistent with a pattern in which high-emission firms earn higher subsequent returns than low-emission firms, particularly over medium- to long-term horizons. This finding aligns closely with the regression-based evidence, which indicates a weakly positive association between carbon intensity and expected returns. In both the regression and portfolio analyses, firms with higher carbon exposure tend to earn higher subsequent returns on average, consistent with carbon-related risks being compensated rather than penalized in equilibrium.

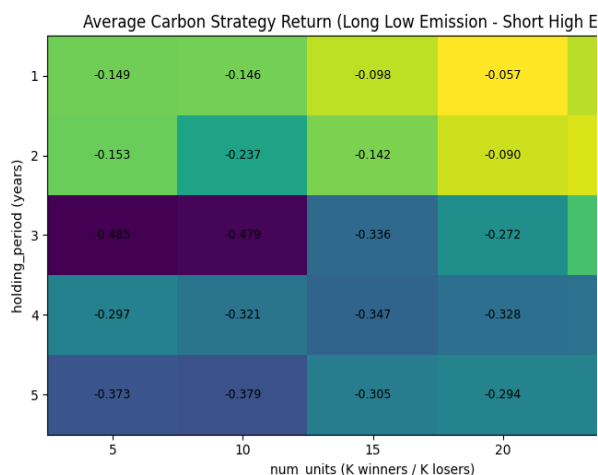


Figure 2: Average Returns of the Carbon-Intensity–Sorted Long–Short Strategy

Notes: The figure reports average net returns from a trading strategy that is long in the K firms with the lowest carbon intensity and short in the K firms with the highest carbon intensity. Rows correspond to holding periods (in years), and columns correspond to portfolio size. Returns are net of proportional transaction costs applied at portfolio entry and exit.

5.6. Carbon-Intensity Decile Strategy

As a complementary exercise to the binary long–short strategy, we implement a carbon-intensity–sorted decile portfolio strategy to examine whether stock returns vary systematically with carbon exposure across the full cross-section of firms. This approach allows for a more granular assessment of return differentials across the full cross-section of firms, rather than focusing solely on the extreme tails.

In each formation year t , firms are sorted into ten deciles based on their lagged carbon intensity, $CI_{i,t-1}$, defined as carbon emissions scaled by sales. Decile 1 (D1) contains firms with the lowest carbon intensity (“green” firms), while Decile 10 (D10) contains firms with the highest carbon intensity (“brown” firms). To ensure robustness of the cross-sectional sort, portfolio formation requires a minimum number of firms in each decile, and years with insufficient cross-sectional coverage are excluded.

Each decile portfolio is held for H years following formation, with $H = 5$ in the baseline simulation. Firm-level holding-period returns are computed by compounding annual simple returns over the holding horizon. Portfolio returns are calculated as equal-weighted averages across firms within each decile. As in the baseline trading simulation, proportional transaction costs are applied at both entry and exit using year-end adjusted prices, and net

portfolio returns are computed after accounting for these fees.

For each formation year, we obtain net returns for all decile portfolios, denoted D1 through D10. To summarize the return differential associated with carbon exposure, we construct a decile spread defined as the return on the low-emission portfolio minus the return on the high-emission portfolio (D1–D10). A positive spread indicates that low-emission firms outperform high-emission firms over the holding horizon, while a negative spread implies relative outperformance of high-emission firms.

The average D1–D10 spread is computed as the time-series mean of holding-period spread returns across all feasible formation years. In addition, the full time series of spread returns and decile-level returns is retained to assess temporal variation in carbon-related return differentials. This decile-based analysis provides a richer characterization of the carbon–return relationship and allows us to evaluate whether returns vary monotonically with carbon intensity, rather than being driven solely by extreme emitters.

The decile strategy results are interpreted as descriptive evidence on the cross-sectional pricing of carbon exposure rather than as an implementable trading strategy. By examining the entire distribution of returns across carbon-intensity deciles, this analysis complements the regression-based evidence and helps assess whether carbon-related return differentials are systematic and monotonic across firms. Together with the baseline long–short strategy, the decile portfolios provide additional insight into how carbon intensity maps into expected returns over medium-term horizons.

Table 4 reports net holding-period returns for portfolios sorted into deciles based on lagged carbon intensity. The decile returns exhibit substantial time-series and cross-sectional variation, with no clear monotonic pattern across all years. In several years, high-emission portfolios (upper deciles) earn relatively high returns, while low-emission portfolios underperform, although this ordering is not uniform across time. The dispersion in decile-level returns highlights the noisy nature of carbon-related return differentials and suggests that carbon exposure interacts with broader market and firm-specific conditions.

The final column of Table 4 reports the return spread between the lowest-emission decile (D1) and the highest-emission decile (D10). The average D1–D10 spread over the sample period is negative, approximately -0.34 , indicating that, on average, low-emission firms underperform high-emission firms over the five-year holding horizon once transaction costs are taken into account. However, the spread varies considerably across years, switching sign in some periods and exhibiting large negative values in others. This volatility underscores that carbon-related return differentials are not stable over time.

Taken together, the decile strategy results suggest that carbon intensity does not generate a monotonic return premium favoring low-emission firms. Instead, higher-emission firms tend to earn higher average returns over medium-term horizons, consistent with compensation for carbon-related transition or regulatory risks. At the same time, the substantial year-to-year variation in the D1–D10 spread cautions against interpreting these patterns as evidence of a stable or easily exploitable trading strategy. Rather, the decile analysis complements the regression-based evidence by showing that carbon exposure is associated with heterogeneous and time-varying return differentials across the full distribution of firms.

The portfolio evidence also helps interpret the regression results in Section 4. While the Bayesian estimates indicate that the average carbon-intensity coefficient is economically small and imprecisely estimated, the carbon-sorted strategies produce sizeable negative low-minus-high spreads. These findings are complementary rather than contradictory for several reasons. First, the regression coefficient captures a conditional marginal effect of carbon

intensity on annual returns after controlling for firm characteristics, market conditions, and hierarchical firm- and sector-level effects, whereas the portfolio sorts reflect largely unconditional return differentials that may embed correlated firm attributes and sectoral composition not explicitly controlled for in the sorting rule. Second, firm-year returns are highly noisy, so posterior uncertainty around a modest average effect is expected in regressions; by aggregating returns across firms and over time, portfolio formation increases the signal-to-noise ratio and can make economically small effects more visible in cross-sectional spreads. Third, the simulations compound returns over multi-year holding horizons, so even a weak annual premium can accumulate into a larger long-horizon differential, consistent with the more negative spreads at longer holding periods. Fourth, the hierarchical model allows for heterogeneous sectoral carbon premia and shrinks noisy sector-level estimates toward the common mean; extreme-sorting strategies may place greater weight on firms and sectors where the carbon premium is relatively stronger, thereby amplifying observed return differences.

Table 4: Carbon-Intensity Decile Portfolio Returns and D1–D10 Spread

Year	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D1–D10
2012	-0.144	0.001	-0.176	–	-0.135	-0.214	-0.313	-0.191	0.359	-0.333	0.189
2013	-0.460	-0.138	-0.248	-0.250	-0.313	-0.477	0.642	-0.049	-0.148	0.528	-0.988
2014	-0.339	-0.271	-0.051	-0.303	0.082	-0.236	-0.121	0.719	-0.425	0.057	-0.396
2015	-0.178	-0.305	-0.069	-0.075	-0.163	0.039	0.253	-0.304	0.429	0.334	-0.513
2016	0.127	0.058	0.712	0.498	-0.031	-0.117	-0.028	0.128	-0.306	0.099	0.028
Average											-0.336**

Notes: This table reports net holding-period returns for carbon-intensity–sorted decile portfolios. Firms are sorted annually into ten deciles based on lagged carbon intensity (emissions scaled by sales), where D1 denotes the lowest-emission firms and D10 the highest-emission firms. Portfolios are held for five years, and returns are equal-weighted and net of proportional transaction costs. The D1–D10 column reports the long–short spread between low- and high-emission portfolios. Years with insufficient cross-sectional coverage to form all decile portfolios are excluded, and missing entries indicate deciles that could not be constructed in a given year. ** denotes statistical significance at the 10% confidence level.

6. Conclusion

This paper examines whether corporate carbon intensity is priced in the cross-section of stock returns, explicitly accounting for sectoral heterogeneity and unobserved firm-level differences. Using a Bayesian hierarchical panel model estimated on Korean firm-level data, we provide an integrated assessment of average, sector-specific, and portfolio-level implications of carbon exposure for expected returns.

Three main findings emerge. First, the average carbon risk premium is positive but economically modest and imprecisely estimated. The posterior distribution of the aggregate carbon-intensity coefficient is centered slightly above zero, yet its highest-density interval includes zero, indicating limited evidence of a strong or uniform carbon

premium across firms once firm characteristics, market conditions, and time effects are controlled for. This result is consistent with recent studies emphasizing that carbon pricing effects tend to be subtle and difficult to identify at the aggregate level (e.g., Bolton & Kacperczyk, 2021; Pankratz et al., 2023).

Second, we find evidence of sectoral heterogeneity in the pricing of carbon exposure. Although individual sector-specific carbon premia are estimated with substantial uncertainty, posterior means are predominantly positive across industries, and the estimated dispersion of industry-level deviations is positive, though modest in magnitude. This pattern suggests that carbon risk is not priced identically across sectors, reflecting differences in production technologies, regulatory exposure, and transition-risk intensity. At the same time, the lack of

precisely estimated sector-level effects cautions against strong conclusions regarding which industries face systematically higher or lower carbon premia. Instead, the evidence points to heterogeneous but noisy pricing of carbon exposure, consistent with the view that climate-related risks are complex, evolving, and only partially incorporated into market prices.

Third, complementary portfolio simulations reinforce the regression-based evidence. Carbon-sorted portfolios indicate that firms with higher carbon intensity tend to earn higher average returns than their low-emission counterparts, particularly over longer holding horizons. Both equal-weighted and value-weighted strategies yield negative returns for low-minus-high carbon portfolios, consistent with investors requiring compensation for bearing carbon-related risks. Importantly, these results do not imply the existence of an exploitable anomaly, but rather align with equilibrium risk compensation in asset pricing, supporting interpretations of carbon exposure as a priced risk factor rather than a mispricing.

Taken together, our findings suggest that carbon-related risks are priced in financial markets in a nuanced and heterogeneous manner. Average effects are small and difficult to estimate precisely, yet cross-industry variation and portfolio-level patterns point to economically meaningful compensation for carbon exposure. These results underscore the importance of explicitly modeling heterogeneity and uncertainty when assessing the financial implications of climate risk and caution against conclusions based solely on average pricing effects.

Several avenues for future research remain. Incorporating forward-looking measures of climate policy risk, firm-level abatement investments, or scope-specific emissions could further illuminate the mechanisms through which carbon exposure affects expected returns. A natural extension is to evaluate carbon-sorted portfolio spreads using richer multifactor benchmarks (e.g., Fama-French style factors constructed for the Korean equity market), which would help separate carbon-related return differentials from exposures to size, value, and other systematic characteristics. In addition, cross-country comparisons across regulatory regimes may help clarify how institutional context shapes the pricing of climate-related risks.

Overall, this study contributes to the climate finance literature by providing an empirical framework that reconciles weak average effects with heterogeneous sectoral pricing and portfolio-level evidence. The results highlight that carbon risk matters for asset prices, but in ways that are subtle, sector-dependent, and inherently uncertain.

Declarations

Ethics Approval and Consent to Participate

This study did not involve human participants or animal subjects.

Competing Interests / Conflicts of Interest

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This work was supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF- NRF-2022S1A5A8054295).

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1), 3–18. [https://doi.org/10.1016/0304-405X\(81\)90018-0](https://doi.org/10.1016/0304-405X(81)90018-0)
- Bates, T. W., Kahle, K. M., & Stulz, R. M. (2009). Why do U.S. firms hold so much more cash than they used to? *Journal of Finance*, 64(5), 1985–2021. <https://doi.org/10.1111/j.1540-6261.2009.01492.x>
- Bhandari, L. C. (1988). Debt/equity ratio and expected common stock returns: Empirical evidence. *Journal of Finance*, 43(2), 507–528. <https://doi.org/10.1111/j.1540-6261.1988.tb03952.x>
- Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2), 517–549. <https://doi.org/10.1016/j.jfineco.2021.05.008>
- Cooper, M. J., Gulen, H., & Schill, M. J. (2008). Asset growth and the cross-section of stock returns. *Journal of Finance*, 63(4), 1609–1651. <https://doi.org/10.1111/j.1540-6261.2008.01370.x>
- Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2), 427–465. <https://doi.org/10.1111/j.1540-6261.1992.tb04398.x>
- Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56. [https://doi.org/10.1016/0304-405X\(93\)90023-5](https://doi.org/10.1016/0304-405X(93)90023-5)
- Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22. <https://doi.org/10.1016/j.jfineco.2014.10.010>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Chapman and Hall/CRC. <https://doi.org/10.1201/b16018>
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.
- Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review*

- of *Financial Studies*, 5(3), 357–386. <https://doi.org/10.1093/rfs/5.3.351>
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623.
- Hong, H., Karolyi, G. A., & Scheinkman, J. A. (2020). Climate finance. *Review of Financial Studies*, 33(3), 1011–1023. <https://doi.org/10.1093/rfs/hhz146>
- Hong, H., Li, F. W., & Xu, J. (2019). Climate risks and market efficiency. *Journal of Econometrics*, 208(1), 265–281. <https://doi.org/10.1016/j.jeconom.2018.09.015>
- Ilhan, E., Krueger, P., Sautner, Z., & Starks, L. T. (2023). Climate risk disclosure and institutional investors. *Review of Financial Studies*, 36(7), 2617–2650. <https://doi.org/10.1093/rfs/hhad002>
- Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1), 65–91. <https://doi.org/10.1111/j.1540-6261.1993.tb04702.x>
- Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708. <https://doi.org/10.2307/1913610>
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1), 1–28. <https://doi.org/10.1016/j.jfineco.2013.01.003>
- Opler, T., Pinkowitz, L., Stulz, R. M., & Williamson, R. (1999). The determinants and implications of corporate cash holdings. *Journal of Financial Economics*, 52(1), 3–46. [https://doi.org/10.1016/S0304-405X\(99\)00003-3](https://doi.org/10.1016/S0304-405X(99)00003-3)
- Organisation for Economic Co-operation and Development (OECD). (2023). *Effective carbon rates 2023: Pricing greenhouse gas emissions through taxes and emissions trading*. OECD Publishing.
- Pankratz, N., Bauer, R., & Derwall, J. (2023). Climate change, firm performance, and investor surprises. *Management Science*, 69(12), 7352–7398. <https://doi.org/10.1287/mnsc.2023.4685>
- Rossi, P. E., Allenby, G. M., & McCulloch, R. (2005). *Bayesian statistics and marketing*. John Wiley & Sons. <https://doi.org/10.1002/0470863692>
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3), 425–442. <https://doi.org/10.1111/j.1540-6261.1964.tb02865.x>
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved \hat{R} for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2), 667–718. <https://doi.org/10.1214/20-BA1221>

variation arises within industries, indicating substantial heterogeneity among firms operating in the same sector. This implies that carbon intensity carries meaningful firm-level information rather than functioning solely as an industry proxy. Consistent with this, our hierarchical model allows for sector-specific deviations in the carbon–return relationship while identifying the common effect primarily from within-industry variation through partial pooling.

4. Because the strategy is constructed as *low-emission minus high-emission* (i.e., long low-carbon, short high-carbon), negative long–short returns imply that high-carbon firms earn higher subsequent returns than low-carbon firms over the holding horizon.

$$5. R_t^{D^1-D^{10}} = \alpha + \beta R_{t \rightarrow t+H}^{mkt} + u_t,$$

6. Numerical annotations are taken from the figure renderings and are approximate. Exact posterior summaries are available upon request.

Endnotes:

1. See Appendix D for additional details on data construction and the emissions-matching process.
2. See Appendix A for the sample coverage diagnostics.
3. To assess whether carbon intensity simply reflects industry structure rather than firm-specific emissions exposure, we decompose its total variation into within- and between-industry components. The decomposition shows that 53.18% of overall

Appendixes

Appendix A: Sample Representativeness and Coverage Diagnostics

This appendix evaluates the representativeness of the emissions-matched analysis sample relative to the full universe of firms listed on the Korea Exchange (KRX). Because firm-level carbon emissions are observed only for entities covered by Korea's mandatory greenhouse gas reporting regime, emissions coverage is incomplete by construction. But it should be noted that emissions-linked firms constitute only a small fraction of listed firms—approximately 10–13% in a typical year over the sample period. At the same time, these firms represent a disproportionately large share of aggregate market capitalization, accounting for roughly 50–60% of total end-of-year market value. This pattern is consistent with the institutional design of the reporting regime, which primarily targets larger and more energy-intensive entities, rather than reflecting data attrition or matching failures. Accordingly, while the analysis sample is limited in terms of the number of firms, it captures a substantial share of the economic weight of the Korean equity market, particularly among firms for which climate-related transition risks are most likely to be material.

To further assess representativeness, Table 5 reports summary statistics for the full KRX firm universe over the sample period, while Table 1 in the main text reports the corresponding statistics for the emissions-matched analysis sample. Comparing these tables provides a transparent diagnostic of how the emissions coverage maps into the broader population of listed firms.

Table 5: Summary statistics for the KRX firm universe

Variable	Mean	Std.	Min	25%	Median	75%	Max
Log assets	20.888	1.841	15.414	19.693	20.844	22.019	26.060
Cash ratio (Cash/Assets)	0.095	0.115	0.000	0.019	0.052	0.118	0.616
Return on assets (ROA)	1.718	9.802	-42.379	0.083	2.121	5.336	30.638
Sales growth (%)	10.799	72.994	-80.186	-9.917	3.346	16.694	387.821
Price-to-book ratio	1.545	1.781	0.094	0.489	0.945	1.811	9.705
Price-to-earnings ratio	24.028	48.742	-219.179	6.180	13.633	25.685	251.912
Market beta	0.920	0.448	-0.133	0.612	0.919	1.227	2.013

Notes: All variables are winsorized at the 1st and 99th percentiles within the sample. Sales growth is reported in percent.

Relative to the full KRX universe, the emissions-matched sample is skewed toward larger firms, as reflected in higher average asset size. Differences in profitability and market beta are relatively modest, while liquidity and valuation ratios show noticeable but economically plausible gaps. These patterns are consistent with the institutional design of Korea's emissions reporting regime, which primarily covers larger, energy- and carbon-intensive firms rather than a random subset of listed companies.

Overall, these diagnostics indicate that the emissions-matched sample differs from the full KRX universe in composition—most notably toward larger firms—while remaining broadly comparable in terms of standard risk characteristics such as market beta. Consequently, the empirical results should be interpreted as describing the pricing of carbon exposure among economically significant, emissions-covered firms rather than the full cross-section of listed firms.

Appendix B. Factor-Adjusted Carbon Portfolio Tests

This appendix provides factor-adjusted analysis of the carbon-intensity–sorted portfolio strategies using CAPM-adjusted spread regressions with heteroskedasticity- and autocorrelation-consistent (HAC) standard errors. By incorporating overlapping holding periods and controlling for market risk, the appendix evaluates whether the return differentials associated with carbon intensity remain robust after accounting for systematic factors.

For each carbon-intensity–sorted long–short strategy, we estimate the following time-series regression:

$$R_t^{LS}(K, H) = \alpha_{K,H} + \beta_{K,H} R_{t \rightarrow t+H}^{mkt} + u_t,$$

where $R_t^{LS}(K, H)$ denotes the net long–short return (low-emission minus high-emission firms) formed in year t with portfolio size K and holding horizon H , and $R_{t \rightarrow t+H}^{mkt}$ is the corresponding H -year compounded market return. The intercept $\alpha_{K,H}$ measures abnormal performance relative to the CAPM, while $\beta_{K,H}$ captures exposure to aggregate market risk.

Standard errors are computed using Newey–West corrections with a lag length of $H - 1$ to account for serial correlation induced by overlapping holding periods. This procedure follows standard practice in the portfolio-evaluation literature and ensures valid statistical inference for long-horizon strategies (Newey & West, 1987; Hodrick, 1992; Jegadeesh & Titman, 1993).

Figure 3 visualizes the estimated CAPM-adjusted alphas $\alpha_{K,H}$ across combinations of portfolio concentration ($K \in \{5,10,15,20,25\}$) and holding horizons ($H \in \{1,2,3,4,5\}$). Each cell reports the estimated alpha along with its HAC-adjusted t -statistic.

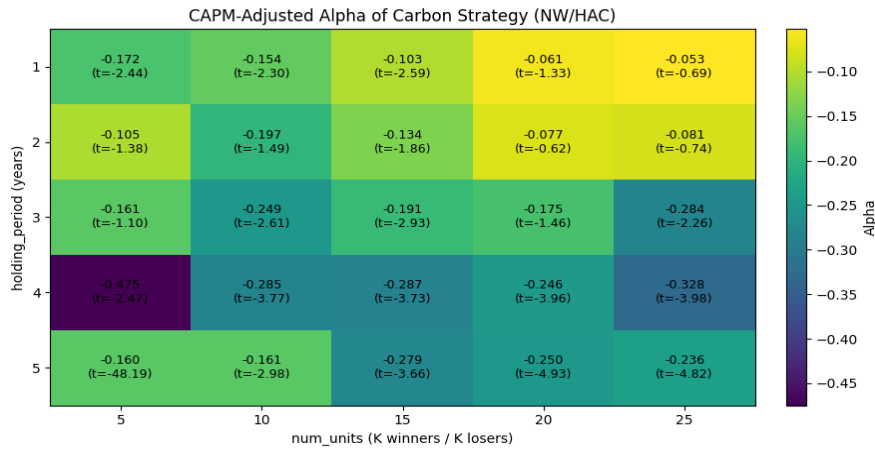


Figure 3: CAPM-Adjusted Alpha of Carbon-Intensity–Sorted Strategies (NW/HAC)

Notes: This figure reports CAPM-adjusted intercepts from regressions of carbon-intensity–sorted long–short portfolio returns on corresponding market returns. Alphas are estimated using Newey–West standard errors with lag length equal to the holding period minus one to correct for overlapping returns. Each cell reports the estimated alpha and its t -statistic in parentheses.

Several clear patterns emerge. First, CAPM-adjusted alphas are uniformly negative across nearly all portfolio configurations, indicating that low-emission firms underperform high-emission firms even after controlling for market risk. Moreover, the magnitude and significance of these negative alphas increase with the holding horizon, with most strategies producing statistically significant alphas for horizons of three years or longer. These findings are stable across portfolio sizes, suggesting they are not driven by a small number of firms.

This factor-adjusted evidence is consistent with the Bayesian regression results in the main text. While the hierarchical model finds a modest and imprecisely estimated average carbon coefficient, sector-level posterior means are predominantly positive. Both approaches point to a weak but positive carbon-related premium that becomes more visible when information is aggregated across firms or over longer horizons.

The CAPM-adjusted portfolio results reinforce this interpretation. Persistent negative alphas for the low-minus-high carbon portfolios imply that high-emission firms earn higher risk-adjusted returns than low-emission firms, and that this differential is not explained by market exposure alone. The spread regressions also yield positive and statistically significant market betas, indicating that carbon-intensive firms load more heavily on aggregate economic risk. This is consistent with carbon intensity capturing exposure to transition risk, regulatory uncertainty, and the cyclicity inherent in energy- and capital-intensive production.

These factor-adjusted findings suggest that carbon intensity contains pricing information beyond simple industry composition or market beta. If carbon exposure merely reflected sector affiliation, CAPM adjustment would eliminate the spread; instead, negative alphas persist—especially at longer horizons—indicating incremental priced variation associated with carbon intensity.

Turning to the decile strategy, CAPM regressions^v for the D1–D10 spread (low minus high carbon) further support this pattern. The estimated alpha is strongly negative ($\alpha = -0.661$, t -value = -6.34), showing substantial underperformance of low-emission firms relative to high-emission firms after accounting for market risk. The positive and significant market beta ($\beta = 1.219$, t -value = 3.19) suggests that high-emission firms are more exposed to systematic macroeconomic fluctuations, consistent with the economic interpretation above.

Together, the negative alpha and positive beta provide a coherent picture: carbon-intensive firms bear greater systematic and transition-related risks and are compensated with higher expected returns. These results align with the Bayesian

hierarchical estimates, which detect only modest average effects at the firm-level but allow sectoral and portfolio-level aggregation to reveal more pronounced patterns.

Overall, the CAPM-adjusted portfolio evidence strengthens the main conclusion of the paper: carbon exposure is associated with higher expected returns, reflecting compensation for transition and regulatory risks. This premium is weak at the average firm-level but becomes clearer in cross-sectional sorting and long-horizon portfolio tests.

Appendix C. Alternative Definitions of Carbon Exposure

This appendix evaluates the robustness of our baseline results to alternative constructions of firm-level carbon exposure. The main analysis (Section 2) defines carbon intensity as verified emissions scaled by sales. We re-estimate the hierarchical Bayesian model (identical priors, Student-t likelihood, identical controls, the same industry random-slope structure, and the same standardization of continuous regressors) under three alternative measures:

1. **Assets-scaled exposure:** $CI_{i,t}^{\text{Assets}} = \text{Emissions}_{i,t} / \text{Assets}_{i,t}$,
2. **Market-cap-scaled exposure:** $CI_{i,t}^{\text{MktCap}} = \text{Emissions}_{i,t} / \text{MktCap}_{i,t}$,
3. **Log-transformed exposure:** $CI_{i,t}^{\text{log}} = \log(1 + \text{Emissions}_{i,t})$.

All alternative measures are lagged one year (i.e., $t - 1$) to align with the information set of investors and to mirror the baseline timing. Each specification replaces the baseline carbon regressor with the corresponding alternative while keeping the covariates, hierarchical components (α_s , δ_i , λ_t), priors, and MCMC settings unchanged. Posteriors are summarized by means and 95% Highest Density Intervals (HDIs).

Figure 4 plots the posterior of the average carbon coefficient when exposure is scaled by assets; Figure 5 scales by market capitalization; and Figure 6 uses the log transform to temper right skew in emissions. Across all alternatives, the posterior distribution of the average carbon coefficient remains small in magnitude and centered slightly above zero, and the 95% HDIs straddle zero.^{vi} Concretely:

- **Assets-scaled:** posterior mean ≈ 0.0052 , 95% HDI $\approx [-0.021, 0.032]$ (Figure 4);
- **Market-cap-scaled:** posterior mean ≈ 0.012 , 95% HDI $\approx [-0.017, 0.042]$ (Figure 5);
- **Log transform:** posterior mean ≈ 0.0096 , 95% HDI $\approx [-0.015, 0.040]$ (Figure 6).

These results mirror the baseline (sales-scaled) finding in Section 4: the average carbon premium is positive but economically modest and imprecisely estimated. Importantly, keeping the identical hierarchical structure implies that identification remains primarily within-industry, and the dispersion of industry-level deviations σ_α is unaffected by the choice of scaling aside from sampling noise. Hence, the conclusion of weak and heterogeneous pricing of carbon exposure is robust to alternative normalizations and to a non-linear (log) transform of emissions.

Scaling by assets and market capitalization addresses potential concerns that sales-based intensity may be correlated with pricing-related fundamentals (e.g., margins, turnover). The log transformation directly targets the heavy right tail in emissions that could otherwise exert undue influence on linear specifications. Under all three alternatives—and with the heavy-tailed Student-t likelihood maintained—the posterior mass remains close to zero with overlapping HDIs, indicating that the lack of a precisely estimated uniform premium is not an artifact of the baseline scaling choice. Together with the portfolio evidence (Section 5), these checks reinforce our interpretation that carbon exposure commands, at best, a small average premium while exhibiting sector-dependent heterogeneity.




Figure 4: Posterior of the Average Carbon Coefficient under Assets-Scaled Exposure

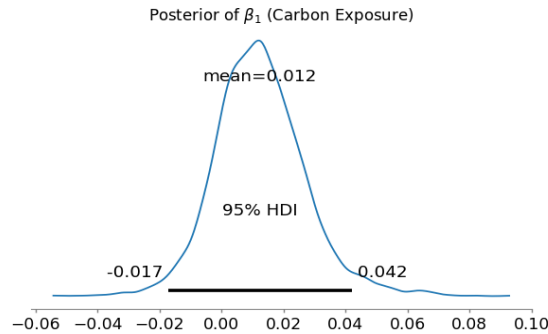


Figure 5: Posterior of the Average Carbon Coefficient under Market-Cap-Scaled Exposure




Figure 6: Posterior of the Average Carbon Coefficient under the Log Transform $\log(1 + \text{Emissions})$

Re-scaling emissions by assets or market capitalization and applying a log transform leaves our core inference intact: the average carbon coefficient is small and statistically uncertain, while the cross-sector pattern remains weakly positive on average. These findings indicate that the paper’s conclusions are not driven by the specific baseline definition of carbon intensity, but persist under economically motivated alternatives.

Table 6 reports sector-specific carbon premia under three alternative exposure measures—emissions scaled by assets, market capitalization, and a log transformation—estimated using the same hierarchical Bayesian model as in the baseline. The table provides a direct comparison to Table 3, which uses sales-scaled intensity. Across all alternatives, the qualitative pattern mirrors the baseline: sectoral posterior means remain mostly positive, especially in carbon-intensive manufacturing and transport sectors, while credible intervals are wide and generally include zero. This indicates a weak but directionally consistent carbon premium with substantial uncertainty.

Although market-cap scaling yields slightly larger posterior means for several heavy-industry sectors, these differences are small relative to overall posterior uncertainty, and sector rankings remain broadly unchanged. The log-transformed measure also produces estimates closely aligned with the baseline, suggesting that results are not driven by outliers or the right-skewed emissions distribution.

Taken together, the evidence shows that sector-level carbon premia are modest and imprecisely estimated regardless of how carbon exposure is defined. This robustness reinforces the conclusion that carbon risk is priced heterogeneously and weakly, rather than through a strong uniform premium, and that the baseline findings are not an artifact of sales-based scaling.

Table 6: Sector-Level Carbon Premia under Alternative Carbon Measures

Sector	Emissions / Assets	Emissions / Market Cap	$\log(1+\text{Emissions})$
Retail Trade (Excluding Motor Vehicles)	0.012 [-0.036, 0.076]	0.021 [-0.040, 0.095]	0.017 [-0.037, 0.086]

Beverage Manufacturing	-0.008 [-0.079, 0.047]	-0.007 [-0.094, 0.059]	-0.004 [-0.078, 0.057]
Land Transport and Pipeline Transport	0.012 [-0.045, 0.088]	0.026 [-0.049, 0.123]	0.014 [-0.043, 0.087]
Professional Services	-0.003 [-0.071, 0.056]	-0.006 [-0.096, 0.059]	-0.000 [-0.073, 0.060]
Motor Vehicles and Trailer Manufacturing	0.012 [-0.026, 0.059]	0.038 [-0.015, 0.111]	0.019 [-0.020, 0.071]
Non-metallic Mineral Products Manufacturing	0.011 [-0.034, 0.071]	0.014 [-0.023, 0.054]	0.022 [-0.034, 0.088]
Electrical Equipment Manufacturing	-0.003 [-0.073, 0.059]	0.003 [-0.080, 0.071]	-0.000 [-0.077, 0.060]
Electronic Components, Computers, Audio, Video and Communication Equipment Manufacturing	-0.010 [-0.084, 0.044]	-0.004 [-0.078, 0.055]	-0.005 [-0.081, 0.048]
Basic Metal Manufacturing	-0.004 [-0.034, 0.025]	0.009 [-0.023, 0.038]	-0.000 [-0.034, 0.028]
General Construction	-0.002 [-0.074, 0.062]	0.005 [-0.076, 0.067]	-0.002 [-0.080, 0.074]
Textile Products Manufacturing (Excluding Apparel)	0.004 [-0.026, 0.032]	0.004 [-0.080, 0.087]	0.011 [-0.020, 0.045]
Pulp, Paper and Paper Products Manufacturing	0.003 [-0.051, 0.053]	0.016 [-0.046, 0.084]	0.007 [-0.050, 0.060]
Food Manufacturing	-0.002 [-0.033, 0.027]	0.013 [-0.024, 0.050]	0.005 [-0.030, 0.037]
Chemical Substances and Chemical Products Manufacturing (Excluding Pharmaceuticals)	0.015 [-0.036, 0.080]	-0.004 [-0.038, 0.026]	0.021 [-0.034, 0.085]
Rubber and Plastics Products Manufacturing	0.014 [-0.043, 0.094]	0.024 [-0.030, 0.100]	0.019 [-0.049, 0.095]
Wood and Cork Products Manufacturing (Excluding Furniture)	0.002 [-0.056, 0.057]	0.023 [-0.054, 0.115]	0.007 [-0.052, 0.070]
Wholesale Trade and Commodity Brokerage	0.003 [-0.062, 0.064]	0.009 [-0.043, 0.062]	0.007 [-0.065, 0.072]
Air Transport	0.005 [-0.058, 0.066]	0.010 [-0.062, 0.096]	0.007 [-0.066, 0.078]
Warehousing and Transport Support Services	0.013 [-0.035, 0.070]	0.006 [-0.073, 0.085]	0.019 [-0.031, 0.077]
Electricity, Gas, Steam and Air Conditioning Supply	0.000 [-0.070, 0.069]	0.023 [-0.043, 0.110]	0.006 [-0.066, 0.073]
Leather, Bags and Footwear Manufacturing	0.004 [-0.057, 0.068]	0.010 [-0.064, 0.101]	0.010 [-0.045, 0.073]
Financial Services	0.005 [-0.064, 0.077]	0.021 [-0.052, 0.119]	0.009 [-0.063, 0.082]
Fabricated Metal Products Manufacturing (Excluding Machinery and Furniture)	0.006 [-0.054, 0.080]	0.016 [-0.060, 0.101]	0.013 [-0.056, 0.083]
Other Transport Equipment Manufacturing	0.002 [-0.061, 0.066]	0.005 [-0.072, 0.083]	0.007 [-0.060, 0.073]
Coke, Briquettes and Refined Petroleum Products Manufacturing	0.015 [-0.030, 0.071]	0.016 [-0.068, 0.101]	0.021 [-0.041, 0.103]
Postal and Telecommunications Services	-0.009 [-0.084, 0.048]	-0.005 [-0.091, 0.064]	-0.007 [-0.092, 0.053]
Computer Programming, Systems Integration and Management Services	0.002 [-0.067, 0.061]	0.014 [-0.064, 0.103]	0.001 [-0.081, 0.072]
Other Machinery and Equipment Manufacturing	0.014 [-0.048, 0.096]	0.036 [-0.037, 0.137]	0.019 [-0.050, 0.105]
Medical, Precision and Optical Instruments and Watches Manufacturing	0.009 [-0.049, 0.079]	0.014 [-0.038, 0.070]	0.014 [-0.048, 0.090]
Insurance and Pension Funding	0.026 [-0.033, 0.120]	0.014 [-0.066, 0.106]	0.034 [-0.029, 0.131]

Tobacco Manufacturing	0.005 [-0.063, 0.067]	0.013 [-0.071, 0.099]	0.010 [-0.059, 0.088]
Sports and Recreation Services	0.013 [-0.055, 0.091]	0.012 [-0.068, 0.103]	0.009 [-0.066, 0.087]
Pharmaceutical Products and Medicinal Materials Manufacturing	0.005 [-0.068, 0.076]	0.012 [-0.067, 0.100]	0.010 [-0.066, 0.090]
Motion Picture, Video and Audio Recording Production and Distribution	0.006 [-0.065, 0.074]	0.012 [-0.077, 0.094]	0.010 [-0.067, 0.083]
Waste Collection, Treatment and Materials Recovery	0.005 [-0.061, 0.079]	0.013 [-0.069, 0.100]	0.010 [-0.063, 0.090]

Notes: Entries report posterior mean and 95% Highest Density Interval (HDI) of the sector-level carbon premium $\beta_c + \alpha_s$ under the hierarchical model, using three alternative exposure definitions. All specifications keep priors, controls, and the Student-t likelihood identical to the baseline; all regressors are standardized. Sectors follow KSIC aggregation used in the main text.

Appendix D. Detailed Data Construction and Matching Procedure

This appendix describes in detail the construction of the analysis dataset, including the matching of emissions records to listed firms and the definition of the variables used in the empirical analysis. Firm-level carbon emissions are obtained from the National Greenhouse Gas Management System (NGMS), which collects verified emissions data from entities subject to Korea's mandatory greenhouse gas reporting requirements. Because these requirements primarily apply to large, energy-intensive establishments, the coverage of emissions data does not extend to the full universe of firms listed on the Korea Exchange (KRX). Consequently, the emissions-linked sample consists of a subset of listed firms that accounts for a disproportionate share of both aggregate emissions and market capitalization.

To integrate these emissions data with financial information, it is necessary to match NGMS reporting entities to publicly listed firms. The NGMS reports verified emissions at the firm level and identifies reporting units using firm names and addresses. To ensure accurate linkage, we rely on firm-level mappings provided by Eco & Partners, an ESG data provider specializing in environmental disclosures. These mappings match listed firms' business registration numbers to the corresponding NGMS reporting entities, allowing CO₂ emission records to be reliably attributed to the appropriate publicly traded firms.

Once emissions have been matched to listed firms, we construct carbon intensity and other firm-level variables used in the empirical analysis. Carbon intensity for firm i in year t is defined as

$$CI_{i,t} = \frac{\text{Emissions}_{i,t}}{\text{Sales}_{i,t}},$$

where emissions are measured in metric tons of CO₂-equivalent and sales are reported in nominal Korean won.

Financial variables are constructed using data from IFRS-based financial statements and market price data. Let $P_{i,t}$ denote the end-of-year split-adjusted stock price of firm i . Annual stock returns are computed as simple returns,

$$R_{i,t} = \frac{P_{i,t} - P_{i,t-1}}{P_{i,t-1}}.$$

Firm size is measured as $\log(\text{Assets}_{i,t})$, where total assets are obtained from IFRS-based financial statements. Leverage is defined as the ratio of total liabilities to total assets at year t . Cash holdings are measured as the cash-to-assets ratio, where cash includes cash and cash equivalents. Profitability is measured by return on assets (ROA),

$$\text{ROA}_{i,t} = \frac{\text{Net Income}_{i,t}}{\text{Assets}_{i,t}},$$

using net income attributable to common equity.

Sales growth is defined as

$$\text{SalesGrowth}_{i,t} = \frac{\text{Sales}_{i,t} - \text{Sales}_{i,t-1}}{\text{Sales}_{i,t-1}},$$

and is expressed in percentage terms in the summary statistics. The price-to-book ratio (PBR) is defined as market capitalization divided by book equity, and the price-to-earnings ratio (PER) is defined as price divided by earnings per share. The contemporaneous market return included as a control variable is the annual simple return on the KOSPI index in year t .

Unless otherwise noted, all regressors are winsorized at the 1st and 99th percentiles within year and standardized to mean zero and unit variance prior to estimation.

¹ See Appendix D for additional details on data construction and the emissions-matching process.

² See Appendix A for the sample coverage diagnostics.

³ To assess whether carbon intensity simply reflects industry structure rather than firm-specific emissions exposure, we decompose its total variation into within- and between-industry components. The decomposition shows that 53.18% of overall variation arises within industries, indicating substantial heterogeneity among firms operating in the same sector. This implies that carbon intensity carries meaningful firm-level information rather than functioning solely as an industry proxy. Consistent with this, our hierarchical model allows for sector-specific deviations in the carbon–return relationship while identifying the common effect primarily from within-industry variation through partial pooling.

⁴ Because the strategy is constructed as *low-emission minus high-emission* (i.e., long low-carbon, short high-carbon), negative long--short returns imply that high-carbon firms earn higher subsequent returns than low-carbon firms over the holding horizon.

^v $R_t^{D1-D10} = \alpha + \beta R_{t \rightarrow t+H}^{mkt} + u_t$,

^{vi} Numerical annotations are taken from the figure renderings and are approximate. Exact posterior summaries are available upon request.