Pricing Transition Risk in European Stock Returns

Master Thesis

Kristian Andersen

Aalborg University
Finance

Finance
Aalborg University Business School
http://www.aau.dk

AALBORG UNIVERSITY

Title:
Pricing Transition Risk in European Stock Returns

Theme:
Asset Pricing of Carbon Transition Risk in European Equities

Project Period:
Spring Semester 2025

Participant(s):
Kristian Andersen

Supervisor(s):
Frederik Lundtofte

Page Numbers: 46

Abstract:

This thesis investigates the pricing of carbon transition risk in European stock returns from 2015 to 2024. We examine whether firm-level carbon exposure, measured by industry-adjusted emissions levels, intensity, growth, and carbon beta (sensitivity to EUA carbon price changes), exhibits a statistically and economically significant relationship with subsequent equity excess returns. Employing quintile portfolio sorts, Fama-French factor models, Fama-MacBeth cross-sectional regressions, and panel fixed-effects models, our unconditional analyses provide limited evidence for a direct carbon risk premium or greenium associated with these metrics. However, a key finding emerges from conditional Fama-MacBeth regressions: the relationship between stock returns and both industry-adjusted emission levels and intensity becomes significantly more positive during periods of rising EUA carbon prices, suggesting state-dependent pricing. This conditional effect, though statistically robust, appears economically modest. Estimated carbon beta was not found to be a distinctly priced risk factor in multivariate settings. Overall, our results indicate that while a simple, unconditional carbon premium is elusive in the recent European market, transition risk pricing is apparent through its interaction with carbon market dynamics

Contents

Chapter 1

Introduction

The global pressure to address climate change is increasingly reshaping economic and financial landscapes, with Europe often at the forefront of policy action. The financial impact of this is underscored by the rapid expansion of carbon pricing mechanisms. A decade ago, such policies covered 7% of global emissions, whereas today, nearly a quarter of global greenhouse gas emissions fall under these instruments, including systems like the EU Emissions Trading System (ETS) (World Bank 2024). This expanding regulatory net creates financial pressures and opportunities for firms, defining the core of "carbon transition risk." Recent data further highlights the effects of this transition, with the European Union, experiencing a significant 15.5% reduction in its ETS-covered emissions in 2023, largely driven by shifts in the power sector (European Comission 2024). Such rapid changes, influenced by policies and in frameworks like the IEA's Stated Policies Scenario (STEPS), emphasize that a firm's ability to navigate this evolving environment is critical for its financial performance (Energy Agency 2024).

In financial economics, a central question is whether investors are compensated for bearing systematic risks (Sharpe 1964; Fama and French 1993). Climate change, and particularly the transition to a low-carbon economy, has emerged as a significant source of such systematic risk (Dietz et al. 2016). Theoretical models, such as Pástor, Stambaugh, and Taylor (2021), present that carbon-intensive ("brown") firms, that is more vulnerable to adverse regulatory or market shifts, should offer higher expected returns i.e. a "carbon premium", to compensate investors. On the other hand, "green" firms might be valued more highly due to investor preferences or their perceived lower risk profile (Pástor, Stambaugh, and Taylor 2021; Krueger, Sautner, and Starks 2020). More people are now checking whether companies' climate plans and statements are truthful, with initiatives like the Science Based Targets initiative (SBTi) and those addressing voluntary carbon market credibility pushing for robust decarbonization efforts beyond simple offsetting (World Bank 2024).

Empirical evidence on a carbon premium has been robust in some markets (Bolton and Kacperczyk 2023; Hsu, Li, and Tsou 2023), with Oestreich and Tsiakas (2015) providing early European insights. However, it's still important to investigate how common carbon risk pricing is and what specific forms it takes, especially in Europe after 2015, given its advanced and fast-changing climate policies. Bolton and Kacperczyk (2023) note that the premium can vary with policy stringency, suggesting that in a region like Europe with active climate policies, the transition risk premia might be nuanced.

This thesis aims to contribute to this discussion by examining the relationship between various measures of carbon exposure and equity returns in the European market from 2015 to 2024. Specifically, this study addresses the following research question:

"For European firms between 2015 and 2024, do measures of carbon transition risk exposure, specifically a firm's emissions level, intensity, growth, or its carbon beta, have a statistically and economically significant relationship with their future excess returns?"

Chapter 2

Theory and Literature

2.1 Theory

2.1.1 Climate Risk in Asset Pricing and the Notion of a "Carbon Premium"

Traditional asset pricing theory holds that investors are rewarded only for bearing systematic (non-diversifiable) risks. In the Capital Asset Pricing Model (CAPM), this means higher expected returns are earned by assets with greater exposure to market-wide risks (Sharpe 1964). Multifactor extensions such as the Fama-French models likewise attribute persistent return differences to risk factors or priced state variables beyond the market (e.g. size, value) (Fama and French 1993). Within this area in finance, climate change has emerged as a new source of systematic risk, and some researchers have begun to ask whether climate-related risks, particularly transition risk, carry a risk premium in equities. Transition risk refers to financial risks from society's transition to a low-carbon economy. For example, policy and regulatory changes, technological disruptions, or shifts in consumer preferences that penalize carbon-intensive business models. If transition risk is systematic, asset pricing theory predicts that investors will demand compensation for bearing it (Pástor, Stambaugh, and Taylor 2021). In other words, firms with greater exposure to transition risk, often measured by higher carbon emissions or intensity, should exhibit higher expected returns i.e. a "carbon premium", all else equal. This idea is somewhat similar to other risk premia: just as firms with high leverage or cyclicality must offer higher returns to compensate investors for added risk, "brown" firms (high emitters) might trade at lower valuations (higher cost of capital) to compensate for the risk of future carbon costs, regulatory shocks, or asset stranding.

Climate risks stand out from traditional factors in important ways. First, climate change manifests through physical risks, e.g damages from extreme weather and transition risks e.g risk from policies, technological advancements and market changes as the world decarbonizes. Physical risks can be sudden disasters or long-term climate shifts, potentially impacting specific regions or industries. Transition risks, on the contrary, are more directly related to carbon emissions and can materialize globally via policy changes (carbon pricing, emissions regulations), technological changes (renewable energy adoption), or changes in investor and consumer behavior. While physical risks have obvious economic impacts, they may sometimes appear idiosyncratic (e.g. a hurricane affecting a region) and thus diversifiable to some extent. Transition risk, on the other hand, is often viewed as a systematic risk factor: for example, a sudden increase in carbon taxes or a technological breakthrough in clean energy could simultaneously reprice assets across multiple high-emission industries. Recent macro-finance assessments suggest climate change can have significant systematic effects on asset values, for instance, Dietz et al. (2016) estimate that unmitigated climate change could shave off a non-trivial fraction of global financial asset values in expectation, with much larger losses in worst-case (tail) scenarios (Dietz et al. 2016). This underscores that climate risk is not only a corporate social responsibility issue but a financial risk factor with economy-wide relevance.

2.1.2 Why Carbon-Intensive Firms May Earn a Risk Premium

From a theoretical standpoint, there are compelling reasons to expect a carbon risk premium. One way this shows up is through risk. Companies that emit a lot of carbon are hit harder by sudden climate rules or shifts in public opinion, so their share prices usually fall when times turn bad, such as when climate policies tighten or climate damage worsens. Investors, being generally risk-averse, dislike assets that crash in bad states and thus will only hold such "brown" stocks if they offer higher expected returns. In an equilibrium model by Pástor, Stambaugh, and Taylor (2021)(Pástor, Stambaugh, and Taylor 2021), this mechanism is formalized: investors dislike unexpected deteriorations in climate, and because carbon-intensive ("brown") firms lose value when climate news turns negative (e.g. new regulations penalizing emissions), those firms are deemed riskier and must offer a higher expected return to attract capital. In their model, climate risk emerges as a state variable alongside the market factor, and brown firms have high "climate betas" (sensitivity to climate risk) (Pástor, Stambaugh, and Taylor 2021). As a result, brown stocks earn positive abnormal returns (CAPM alphas) in equilibrium as compensation for this additional systematic risk exposure. In contrast, "green" assets (low-emission firms) tend to hedge climate risk, they fare relatively better when climate-related shocks hit, so they provide insurance-like properties to investors. Much like other hedging assets, green stocks are predicted to have lower expected returns, even yielding negative alphas relative to CAPM, because investors are willing to accept lower returns in exchange for their climate-hedging benefits (Pástor, Stambaugh, and Taylor 2021). This theoretical insight aligns with the classic risk–return tradeoff and implies a persistent "brown premium": higher carbon intensity is associated with higher required returns, all else equal.

A second channel comes from investor preferences and constraints. A growing segment of investors exhibit environmental, social, and governance (ESG) preferences, meaning they derive utility from holding "green" assets or disutility from holding "brown" ones. In recent decades, trillions of dollars have flowed into sustainable or low-carbon investment strategies (Krueger, Sautner, and Starks 2020). These preference-driven flows can directly affect asset prices: if many investors avoid or divest from carbon-intensive firms ("exclusionary screening"), those stocks will trade at a discount i.e. lower demand, thereby raising their expected returns. Pástor et al. (2021) (Pástor, Stambaugh, and Taylor 2021) incorporate such taste-based investing into their equilibrium model and show that even absent any true risk difference, green assets can have lower expected returns simply because investors enjoy holding them. In reality, both channels likely operate simultaneously: some investors demand a return premium for carbon risk, while others are willing to sacrifice returns to tilt towards greener assets. The net effect in market prices is subtle. If climate-concerned investors are significant enough, they can drive down green firms' cost of capital and inflate brown firms' cost of capital (i.e. brown stocks become cheap, with high future returns, whereas green stocks become expensive, with lower future returns). Notably, this does not violate economic logic, it is akin to "sin" stocks (e.g. tobacco, alcohol) historically earning abnormal returns because they were shunned by certain investors (Hong and Kacperczyk 2009). Climate-conscious investing can therefore lead to a "greenium" (green assets priced at a premium, yielding lower returns) and a corresponding brown premium, consistent with equilibrium outcomes where green assets underperform over the long run (Pástor, Stambaugh, and Taylor 2021). In sum, whether due to risk or tastes, theory suggests carbon-intensive firms could indeed face a higher cost of capital and offer higher expected returns than their low-carbon peers.

It is important to clarify that a carbon risk premium does not imply that high-carbon stocks will always realize higher returns in every sample period, rather it means they promise higher returns ex ante as compensation for risk. If a severe transition shock occurs, e.g. a harsh carbon tax or rapid technological shift away from fossil fuels, brown firms may suffer large losses in that scenario. Investors who held those firms would then incur the downside of the risk they were being compensated for. Thus, over short horizons or specific periods of realized climate policy shocks, green firms can actually outperform (since the risk has materialized). Empirical studies document that in news-sensitive periods green stocks behave like "crisis hedges" – for example, during abnormally warm months that heighten awareness of climate change, low-emission (green) firms significantly outperform high-emission firms. Choi et al. (2020) (Choi, Gao, and Jiang 2020) show that extreme heat waves, which likely remind investors of climate risks, trigger relative selling of carbon-intensive stocks and flight to greener stocks. Similarly, Engle et al. (2019) (Iii et al. 2019) find that stocks of companies with strong environmental scores outperform in periods of negative climate news . These patterns imply that green portfolios act as hedges when climate concerns spike, consistent with the notion that brown stocks carry the downside risk. Over the long run, however, it is precisely this downside exposure that justifies higher equilibrium returns for brown firms. Thus, short-term episodes of green outperformance are not inconsistent with a longer-term brown risk premium; rather, they are the manifestation of the risk materializing in those moments. This dynamic underscores why the "transition risk premium" is a complex phenomenon – it intertwines risk compensation with shifting investor sentiment and policy expectations over time.

2.1.3 Empirical Evidence on Carbon Intensity and Expected Returns

A growing empirical literature has tested whether carbon emissions (a proxy for transition risk exposure) predict stock returns. Early studies provided intriguing evidence, particularly in Europe. Oestreich and Tsiakas (2015) (Oestreich and Tsiakas 2015) examined stock performance during the initial phase of the EU Emissions Trading Scheme (EU ETS), the world's first major carbon cap-and-trade program, and found that German companies receiving large free allocations of carbon permits significantly outperformed those that received none. They interpret this as a "carbon premium" in equity returns during the early EU ETS years. Part of this premium came from the immediate cash-flow boost firms received when they were granted free carbon allowances, which in turn lifted their share prices. But importantly the authors also identify a risk-based component, where firms with higher carbon emissions exhibited higher exposures to a carbon price factor and tended to have higher expected returns than cleaner firms (Oestreich and Tsiakas 2015). In other words, even in a highly regulated market like the EU, carbon-intensive firms had a higher cost of capital consistent with investors pricing carbon risk. This European evidence was among the first to link carbon emissions to cross-sectional return differences.

Subsequent research generalized these findings across broader markets. A landmark study by Bolton and Kacperczyk (2021) (Bolton and Kacperczyk 2021) analyzed U.S. stock returns and documented that firms with higher CO_2 emissions earn higher subsequent returns, even after controlling for size, value, momentum, and other known factors. This carbon premium was statistically significant and could not be explained by differences in profitability or other firm characteristics. In fact, the authors remark that "we cannot explain this carbon premium through differences in unexpected profitability or other known risk factors" (Bolton and Kacperczyk 2021), which suggests that the premium is not just explained by other known factors or a short-lived anomaly. Additional analyses in their study showed that certain institutional investors were actively avoiding high-carbon stocks, for example, pension funds or funds with ESG mandates tended to underweight "brown" firms and as a result, these investors sacrificed some returns by doing so. This behavior is consistent with an equilibrium where carbon risk is priced: investors "care" about carbon risk enough to tilt portfolios, and those who bear the carbon exposure are rewarded with extra return (Bolton and Kacperczyk 2021). In follow-up work, Bolton and Kacperczyk (2023) (Bolton and Kacperczyk 2023) expanded the analysis globally, examining 14,400 firms across 77 countries. They again found a robust carbon risk premium: stocks of companies with higher carbon emission levels and higher emission growth rates delivered higher returns in most sectors and countries. Notably, the magnitude of the premium varied with the economic and policy environment. For instance, the return premium associated with emission levels was larger in countries with more stringent climate policies. at first glance, this is puzzling as one might expect strong climate policy to punish brown firms' performance. The interpretation offered is that in countries with aggressive climate regulation (such as many in Europe), carbon-intensive firms are perceived as especially risky (since they face greater transition shocks), and thus they trade at steeper discounts, i.e. higher expected returns, to compensate investors (Bolton and Kacperczyk 2023). In contrast, in countries with lax climate policy, investors might not price carbon risk as strongly. Likewise, Bolton and Kacperczyk (2023) (Bolton and Kacperczyk 2023) find that the premium related to emission changes is higher in emerging markets and countries with less inclusive governance, perhaps reflecting that uncontrolled emissions growth in such environments is viewed as a sign of unmanaged risk, requiring higher returns. The overarching conclusion from these studies is that a systematic "carbon premium" has been present: high emitters have delivered excess returns relative to low emitters, consistent with investors demanding a premium for transition risk exposure.

Importantly, these return patterns have persisted even when accounting for the classic Fama-French factors and industry effects, suggesting that carbon intensity is capturing a distinct dimension of risk. For example, many carbon-intensive firms are in sectors like energy, utilities, or materials, which often have value stock characteristics (high book-to-market). Yet the carbon-related return spread does not simply reduce to the value premium or industry performance, it appears as an independent factor. A long-short portfolio that is long "brown" (high emission intensity) stocks and short "green" (low emission) stocks within the same industry tends to earn positive abnormal returns in both U.S. and European samples. Hsu, Li, and Tsou (2023) (Hsu, Li, and Tsou 2023) construct such a portfolio based on firms' toxic emission intensities (a proxy for pollution and carbon intensity) and find an average outperformance of about 4.4% per year for the brown-minus-green strategy. This "pollution premium" remains significant even after controlling for the market, size, value, momentum, and other standard factors. In their analysis, this premium cannot be explained away by firm characteristics like profitability, investment, or even by differences in investor sentiment or governance – pointing again to an underlying risk factor. The authors propose that the premium reflects systematic regulatory risk: specifically, they model a factor related to environmental policy uncertainty and show that firms with high pollution exposure load on this factor (Hsu, Li, and Tsou 2023). When the risk of stricter environmental regulation rises, these firms' valuations suffer and investors demand higher returns as compensation for this ever-present threat of a regime change. Hsu et al. (2023) (Hsu, Li, and Tsou 2023) provide direct evidence by linking the return spread to a proxy for regulatory risk (growth in environmental litigation and penalties) and find that this measure helps explain the cross-section of returns on emission-sorted portfolios. This aligns neatly with the notion of a transition risk premium: firms more exposed to future carbon costs are systematically riskier and priced accordingly.

Another complementary strand of evidence comes from option markets and tail risk. Ilhan, Sautner, and Vilkov (2021) (Ilhan, Sautner, and Vilkov 2021) examine options on U.S. stocks and find that investors are willing to pay more to insure against downside risk in carbon-intensive firms. Specifically, firms with higher carbon intensity have steeper implied volatility skews (more expensive put options relative to calls), indicating a perceived higher tail risk for those stocks. Crucially, this "carbon tail risk" premium in option prices becomes more pronounced when climate news is negative (Ilhan, Sautner, and Vilkov 2021). The authors also document a natural experiment: after the surprise 2016 U.S. election of a climate-policy-skeptic president (an event that reduced expected transition risk), the option-implied tail risk for high-carbon firms subsided significantly (Ilhan, Sautner, and Vilkov 2021). This suggests that part of what investors were insuring against, a sudden transition shock, became less likely, and the price of that insurance, the tail risk premium, fell. Such findings reinforce that transition risk is recognized and priced in markets: equity investors demand compensation ex ante (higher returns, lower valuations) and even buy protection (options) to hedge against the worst outcomes for carbon-intensive firms. In a sense, the options market is revealing the same story as the stock returns: carbon-heavy companies carry an extra downside risk that investors price in.

While the weight of evidence supports the existence of a carbon-related risk premium, the literature is not entirely uniform, and several unresolved questions remain. Some studies find that once certain adjustments are made, the carbon premium can appear attenuated or even absent. For example, Görgen et al. (2020) (Riordan n.d.) construct a carbon risk factor and find that although brown firms had higher average returns historically, a portfolio tracking carbon-risk did not earn a statistically significant premium in their sample. They observe that firms which improved their carbon footprint ("became greener") tended to get a positive price reaction (lower expected returns going forward), offsetting some of the static brown-vs-green return difference. Moreover, they argue that a lot of the carbon exposure corresponded to cash-flow news rather than discount-rate news – in other words, high emitters experienced stock price drops when, say, emissions increased i.e. bad news about future cash flows due to potential regulation, but these price moves were not necessarily compensation for risk in the usual sense (Riordan n.d.). As a result, in their tests the carbon factor did not carry a significant risk premium once these effects were accounted for. This perspective raises a subtle point: is the observed "carbon premium" truly a reward for bearing risk, or partly a result of mispricing or other frictions? If some investors systematically avoid brown stocks, those stocks could become undervalued and earn higher returns temporarily until corrected – not because of risk per se, but because of a persistent demand shortfall (a pricing effect rather than a risk effect). Distinguishing between these explanations is challenging and is an ongoing debate.

2.1.4 Summary

In summary, the literature provides a strong conceptual and empirical basis for the idea that carbon transition risk is priced in equity markets. Classic asset pricing frameworks can accommodate a climate risk factor, and equilibrium models predict a systematic return differential between brown and green firms as a function of both risk exposure and investor tastes (Pástor, Stambaugh, and Taylor 2021). Empirical studies, especially in the U.S. and global samples, have overwhelmingly found that high-emission firms earn higher average returns than low-emission firms, consistent with a carbon risk premium (Bolton and Kacperczyk 2021; Bolton and Kacperczyk 2023; Hsu, Li, and Tsou 2023). This pattern holds after accounting for traditional factors, and seems linked to the risk of future climate policy or demand shifts that weigh on brown firms. At the same time, there is evidence that investors are increasingly attentive to climate issues, in some cases willingly forgoing returns to hold greener portfolios, which complicates the interpretation of the premium (Bolton and Kacperczyk 2021). The distinction between transition risk (e.g. policy-driven revaluation risk) and physical risk is also important: most of the return premia observed relate to transition (emissions exposure) rather than physical climate damage, which suggests that carbon intensity is serving as a proxy for regulatory and technological disruption risk. Physical climate risks may play out over longer horizons and could be more idiosyncratic at the firm level, making them harder to detect in cross-sectional returns. Though certain sectors like insurance or real estate do show sensitivity to physical climate events, as shown in Chava et al. (2014) or Dietz et al. (2016). Thus, the focus in asset pricing has understandably been on emissions and transition risk as the more immediate systematic factor.

2.1.5 Research Gap

Unresolved questions remain regarding how stable and pervasive the carbon premium is, particularly in different markets and time periods. Europe represents a particularly interesting case. European equity markets have been at the forefront of climate policy (with mechanisms like the EU ETS and stronger regulatory mandates) and European investors are often cited as among the most climate-conscious (Krueger, Sautner, and Starks 2020). Does this mean the carbon premium is higher in Europe (because transition risk is very salient), or could it be lower/absent (if European markets have already priced in climate risks more efficiently, or if investor demand for green assets is especially strong)? The empirical evidence specific to Europe is still relatively sparse compared to U.S. and global studies. Apart from the early EU ETS study (Oestreich and Tsiakas 2015) and a few multi-country analyses, there is room to clarify how carbon risk is priced across European equities in the post-2010 period, especially as EU climate policies have tightened (e.g. the ramping up of emissions targets for 2030 and beyond). Moreover, prior studies have used various measures of emissions, levels (total CO_2), intensities (emissions per revenue or asset), and changes or growth in emissions, as predictors. Each measure has a slightly different interpretation: intensity controls for firm size and efficiency, while changes capture improvement or deterioration in carbon footprint. Bolton and Kacperczyk's (2023) (Bolton and Kacperczyk 2023) global evidence suggests both matter, but it is not fully answered which aspect is most strongly rewarded in Europe's context. This thesis aims to contribute new evidence on these issues. In the following chapters, we develop an empirical asset pricing analysis for European equities, testing whether carbon emission intensity (our proxy for transition risk) commands a risk premium. We will build on the cited theoretical frameworks, treating carbon intensity as a potential risk factor, and on the empirical methodologies of prior studies (portfolio sorts, factor regressions, etc.) to examine returns. By focusing on European markets, we also address a geographical gap and explore whether the patterns observed in the U.S. and global data hold under Europe's regulatory and investor environment. Furthermore, our analysis will explore related questions such as whether changes in emissions (improvements or worsening) carry return implications (i.e. are investors rewarding firms that decarbonize, or is the premium primarily attached to static high-emission exposure?). Ultimately, this theoretical background underlines why carbon transition risk might be priced and how that pricing is observed in practice as a return premium for high-carbon firms. The empirical investigation that follows will shed further light on these relationships, helping to clarify whether the transition to a low-carbon economy is reflected in the cost of equity capital for European companies, and by extension, how investors are, or are not, pricing the risks of climate change in financial markets.

Chapter 3

Methodology

3.1 Methodology

3.1.1 Data Description

We employ a panel of European equities from the STOXX600 index from January 2015 to December 2024, covering firms across various industries, classified by ICB supersectors. Monthly stock returns for each firm are used as the primary dependent variable. yearly firm fundamentals is gathered such as market capitalization (size), book-to-market ratio, return on equity, and recent stock performance (momentum). To mitigate the impact of outliers, key continuous variables (returns, B/M, etc.) are winsorized at the 1% level. All monetary values are in euros.

Carbon emissions data Scopes 1 and 2 are obtained from Refinitv Eikon . Scope 1 and 2 emissions represent direct operational carbon output and purchased energy emissions (McKinsey 2024). Because emissions are reported annually with a lag, we align emissions data with returns carefully. Emissions figures are lagged by one year in the analysis to avoid look-ahead bias and reflect the information available to investors. This means, for example, that a firm's emissions from fiscal year 2020 are used to explain stock returns from 2021. Using lagged emissions is essential since carbon data are typically released many months after year-end. We also preprocess emissions by taking natural logarithms or scaling them (as detailed below) to reduce skewness given the heavy-tailed distribution of emissions across firms.

In addition to firm-level data, we incorporate two key climate risk variables at the market level. First, the European Union Allowance futures price series is included to capture carbon market shocks. We use the monthly change in EUA prices (ΔEUA) as a proxy for surprises in carbon cost where an increase in EUA price represents a tightening of carbon regulation or higher carbon cost for emitters. Second, we include a Climate Policy Uncertainty index, measured monthly, which reflects the level of uncertainty in climate-related policy and regulation. The CPU index is constructed from news-based economic policy uncertainty indicators focusing on climate keywords (Gavriilidis n.d.). By including ΔEUA and the CPU index, we control for time-varying macro-level forces that could affect all high- or low-carbon firms simultaneously e.g. new climate legislation or carbon tax news. These series serve as additional risk factors and control variables in our regressions, ensuring that our results on carbon risk pricing are not disturbed by broad market trends in carbon pricing or policy uncertainty.

3.1.2 Carbon Measures

A core metric in the methodology is the construction of multiple emissions metrics to capture different facets of a firm's carbon exposure. We define and use the following measures:

Absolute Emissions (Levels)

The total greenhouse gas emissions of the firm. We consider Scope 1+2 emissions as one measure of operational carbon footprint. Using absolute emissions reflects the firm's total contribution to carbon output. Prior literature suggests that firms with higher absolute emissions have tended to earn higher returns, consistent with a carbon risk premium (Bolton and Kacperczyk 2023).

Carbon Intensity (Emissions/Revenue)

Following standard practice, we scale emissions by firm size to obtain an emissions intensity metric. We define carbon intensity as total emissions divided by the firm's annual revenue i.e tons of CO_2 per euro of revenue. This normalizes emissions by economic output, thus distinguishing a firm's environmental impact from its size. Intensity measures a firm's carbon efficiency; high values indicate a firm emits a lot of CO_2 for each unit of revenue, making it dirtier relative to its economic activity. The literature argue that intensity is often a more informative measure than absolute emissions, since absolute emissions tend to scale roughly one-for-one with firm sales (Ilhan, Sautner, and Vilkov 2021). By using intensity, we control for the fact that larger firms naturally emit more, and we isolate the carbon footprint relative to operations. In computing intensity, we lag the emissions and use revenue from the same fiscal year as the emissions figure.

Industry-Adjusted Carbon Intensity

To further isolate firm-specific carbon impact, we compute an industry-adjusted intensity z-score for each firm-year. This is done by subtracting the industry-year mean intensity from the firm's intensity and dividing by the industry-year standard deviation. Essentially, we standardize each firm's carbon intensity relative to peers in the same industry and year. This yields a within-industry measure of carbon intensity, highlighting how much higher or lower a firm's emissions are relative to companies with similar business activities. The argument is that industries differ greatly in emissions (utilities vs. software, for example), so an industry-adjusted metric captures the firm's carbon performance after controlling for industry-wide factors. For example, a coal power company and a renewable energy company might both be utilities, but the coal plant would have a very high industry-adjusted intensity. We use this standardized score in industry-neutral portfolio sorts to ensure our results are not driven purely by sector composition.

Emissions Growth

We also consider changes in emissions over time as an indicator of transition trajectory. Emissions growth is measured as the year-over-year log difference in emissions (e.g. $\log(\text{Emissions}_t) - \log(\text{Emissions}_{t-1})$). We compute this for Carbon Intensity (and in some cases separately for Scope 3). This metric captures how quickly a firm is decarbonizing (or increasing its emissions). A large negative value indicates a firm that sharply cut emissions, whereas a large positive value indicates rapidly rising emissions. Prior research has examined changes alongside levels, finding that both higher levels and increases in emissions can predict higher returns (Bolton and Kacperczyk 2023). Including emission growth allows us to test whether markets price not just the static dirtiness of a firm but also its direction of change (improvement or escalation in carbon footprint).

Each of these measures serves a purpose. Absolute emissions measures total exposure to carbon costs. Carbon Intensity measures carbon efficiency and relative performance on climate metrics, which may align with investor preferences for "greener" companies or with risk exposure to future regulations. By examining both levels and intensities, we acknowledge the ongoing debate in the literature about which aspect of emissions is most relevant for asset pricing. Our approach is inclusive: we expect that if carbon transition risk is priced, it could manifest through one or several of these metrics. All emissions data are winsorized at extreme percentiles to avoid spuriously large ratios or growth rates driving results.

3.1.3 Portfolio Construction

To investigate the carbon transition risk premium, we construct portfolios sorted on firms' carbon characteristics and examine their performance. Specifically, we form quintile portfolios based on carbon intensity. At the end of each month (assuming updated emissions data are available annually, we keep the sorting criterion constant between updates), we rank all stocks by their carbon measure and assign them into five quintiles. Portfolio 1 (Green) contains the 20% of firms with the lowest carbon intensity (i.e. "green" firms), whereas Portfolio 5 (Brown) contains the 20% with the highest intensity (the "brown" firms). Each portfolio's return in the subsequent month is calculated as a value-weighted average of the constituent stocks' returns, using each stock's market capitalization at the previous month-end as weights. Value-weighting ensures that the portfolio performance is representative of an investable strategy and not unduly influenced by tiny firms. We rebalance these portfolios monthly, allowing membership to change as firms' updated carbon data (lagged one year) or market caps evolve.

From these deciles, we derive a factor-mimicking portfolio for carbon risk. The "Brown Minus Green" (BMG) portfolio is a zero-investment long–short strategy that goes long the highest-emission quintile and short the lowest-emission quintile each month. This carbon factor represents the return spread between carbon-intensive and carbon-efficient firms. Formally, it is the return difference between Portfolio 5 and Portfolio 1 (brown minus green). This construction is similar to approaches in recent literature that create long–short factors for ESG or carbon risk. In essence, $BMG_t = \bar{R}_{\text{Brown},t} - \bar{R}_{\text{Green},t}$, which yields a time series of the carbon transition premium. A positive average return to BMG would indicate that brown firms outperform green firms on average (consistent with investors demanding a premium to hold carbon-intensive stocks), whereas a negative average would imply green firms outperform (perhaps due to investor preferences or underestimated growth prospects). We compute the mean and volatility of this factor and track its behavior over time.

We also construct an industry-neutral version of the carbon factor to ensure our results are not purely driven by sectoral differences. To do this, we use the industry-adjusted intensity metric for sorting. Each month, we rank stocks by their industry-relative carbon intensity and form quintile portfolios similarly, then take a long-short High-Low position. This produces an industry-adjusted BMG factor where, by construction, the high- and low-carbon portfolios have a similar industry composition (each industry's representation in the long and short sides is approximately equal). Alternatively, one can think of this as forming long–short spreads within each industry and averaging them: effectively hedging out broad industry effects. The industry-neutral carbon factor thus captures the return difference between firms that are unusually carbon-intensive for their industry and those that are relatively green within the same industry. This helps determine if any carbon premium is truly a within-sector phenomenon (e.g. dirty vs. clean power companies) rather than just the result of heavy-emitting industries (energy, materials) having different performance than low-emitting industries (tech, services).

All portfolios are constructed using excess returns (return minus the risk-free rate) for consistency with factor modeling. We ensure that the long–short portfolios are "dollar-neutral" each month. We also examine the characteristics of these portfolios to verify that our green vs. brown classification aligns with intuitive differences (for example, the brown quintile should have higher average emissions and may be overweight utilities, materials, etc., unless neutralized by industry adjustment). The time series of the BMG factor will be used in subsequent asset pricing tests to see if it contains unique information not captured by standard factors.

3.1.4 Factor Model Analysis

We begin by evaluating the carbon factor in a time-series context alongside well-known risk factors. In particular, we test whether the BMG (brown-minus-green) factor's returns can be explained by existing asset pricing models or if it delivers a non-zero alpha (abnormal return). To do this, we run time-series regressions of the carbon factor on the Fama–French factors and other controls. Specifically, we estimate regressions of the form:

$$R_{BMG,t} = \alpha + b_{\text{MKT}}\, R_{\text{MKT},t} + b_{\text{SMB}}\, SMB_t + b_{\text{HML}}\, HML_t + b_{\text{RMW}}\, RMW_t \\ + b_{\text{CMA}}\, CMA_t + c_1\, \Delta\text{EUA}_t + c_2\, \text{CPU}_t + \varepsilon_t \tag{3.1}$$

where $R_{BMG,t}$ is the excess return of the carbon long–short portfolio in month t. We consider both the Fama–French 3-factor model (market, size, value) and the 5-factor model (adding profitability and investment factors) as baselines. The coefficients measure the exposure of the carbon factor to each systematic factor. In addition, we include ΔEUA (monthly change in carbon price) and CPU (climate policy uncertainty) as additional regressors to see if the carbon factor is correlated with innovations in climate policy or carbon pricing beyond broad market movements. These can be viewed as non-traditional factors that might drive the carbon spread.

From these regressions, we focus on the alpha (α) and the factor loadings on standard factors. A statistically significant α would imply that the carbon factor earns excess returns that are not accounted for by exposures to known risk factors, indicating the presence of a distinct carbon risk premium. We report the alpha in monthly percentage terms and its t-statistic. We also examine the sign and magnitude of loadings. Perhaps, we might find that BMG has a negative loading on the market factor if green stocks have lower betas on average or vice versa. We expect, before looking at the data, that the carbon factor could be somewhat correlated with value (HML) or profitability (RMW) since high emitters might be asset-heavy "value" firms, but these relationships are empirical questions. We report adjusted R^2 to assess how much of the carbon factor's variation is explained by the model. If adding our carbon factor significantly improves the explanatory power for other portfolio returns, that would indicate it carries incremental information, but here we are specifically checking the converse – how well existing factors explain the carbon factor. Standard errors in these time-series regressions are adjusted using Newey–West corrections (lags chosen based on monthly data frequency, 6 lags) to account for any autocorrelation in factor returns.

In summary, the factor model analysis tests whether the "brown-minus-green" return spread is an independent source of risk. A finding of zero alpha would suggest that the carbon spread is fully attributable to known factors for example, if brown firms are just value stocks in disguise, whereas a positive alpha would imply that investors earn additional returns for bearing carbon transition risk. I also interpret the coefficients on CPU and ΔEUA: for instance, a positive coefficient on ΔEUA in this regression would mean the carbon factor tends to be higher (brown outperforms green) in months when carbon prices rise, or negative would mean green outperforms when carbon prices jump. These patterns help us understand the nature of the carbon factor (e.g., whether brown stocks are particularly sensitive to carbon price shocks).

3.1.5 Firm-Level Regressions

To complement the portfolio-level analysis, we conduct firm-level return regressions to examine whether carbon exposure is priced in the cross-section of individual stocks. We use two main approaches: Fama–MacBeth cross-sectional regressions and panel fixed-effects regressions. These techniques allow us to control for firm characteristics and time effects while assessing the marginal effect of carbon metrics on returns.

Cross-Sectional Regressions (Fama–MacBeth)

We follow the two-step Fama–MacBeth (1973) procedure to estimate risk premia associated with carbon metrics. In the first step, for each month t in our sample, we run a cross-sectional regression of individual stock excess returns (monthly) on firm characteristics measured at the end of the previous period. The regression takes the form:

$$\begin{aligned} R_{i,t} - R_{f,t} &= \gamma_{0,t} + \gamma_{1,t}\,\mathrm{CarbonMetric}_{i,t-1} + \gamma_{2,t}\,\ln\big(\mathrm{Size}_{i,t-1}\big) \\ &\quad + \gamma_{3,t}\,\ln\Big(\tfrac{B_{i,t-1}}{M_{i,t-1}}\Big) + \gamma_{4,t}\,\mathrm{ROE}_{i,t-1} + \gamma_{5,t}\,\mathrm{Momentum}_{i,t-1} + \varepsilon_{i,t}. \end{aligned} \tag{3.2}$$

Here, CarbonMetric represents one of our emissions measures for firm i (e.g. emissions growth or industry-adjusted intensity, lagged to ensure it's known at time t). We include Size (log market cap), Book-to-Market (log of book value-to-market value), Profitability (return on equity), and Momentum (cumulative stock return over the past 12 months excluding the last month) as control variables, as is standard in cross-sectional asset pricing tests. These controls account for the well-established Fama–French factors and other anomalies, ensuring that any relation between carbon and returns is not due to omitted common factors. Each month we obtain a vector of slope estimates $\gamma_{j,t}$. In the second step, we compute the time-series average of each coefficient and assess its significance using Newey–West adjusted standard errors (to correct for serial correlation in monthly estimates). This yields an estimate of the price of risk associated with each characteristic. For example, a significantly positive coefficient on carbon intensity would indicate that, on average, high-carbon firms earned higher subsequent returns than low-carbon firms (even after controlling for size, value, etc.), consistent with a carbon risk premium. We report these average coefficients and t-statistics. Key continuous independent variables used in these regressions are globally winsorized to reduce the influence of extreme observations. The Fama–MacBeth approach is attractive here because it directly tests cross-sectional relationships in each period and produces standard errors that account for the time-series variability of these relations.

Panel Fixed-Effects Regressions

While Fama-MacBeth focuses on cross-sectional average effects, we also exploit the panel structure of our data to run firm-level panel regressions with fixed effects. The general model is:

$$R_{i,t} - R_{f,t} = \beta\,\text{CarbonMetric}_{i,t-1} + X_{i,t-1}\delta + \mu_i + \lambda_t + \varepsilon_{i,t} \tag{3.3}$$

where μ_i are firm fixed effects and λ_t are time (month) fixed effects, and $X_{i,t-1}$ includes the same firm characteristics as above (size, B/M, etc., lagged). The firm FE absorb any time-invariant differences between companies (such as industry affiliation or persistent governance traits), so the carbon coefficient β is identified from within-firm variation over time. The time FE absorb any month-specific shocks common to all firms (market-wide returns, macro news). This setup means we are effectively asking: when a given firm's carbon intensity is higher or lower than its own average (or when it increases its emissions relative to before), do its returns tend to be higher or lower, after accounting for market-wide movements and other traits? We estimate β using the entire panel (2015–2024) and cluster standard errors by both firm and month to allow for arbitrary correlation in residuals across firms in a given time and within a firm over time (two-way clustering addresses both cross-sectional and time-series dependence) – a robust approach recommended for panel data. A significant β in this fixed-effects model would indicate that carbon intensity has a predictive effect on returns even when comparing a firm to itself over time and netting out broad market trends. I note that firm-level carbon metrics usually change slowly year to year, so the power of this within-firm test may be lower; we include it mainly to ensure that any cross-sectional findings are not driven by omitted firm-specific factors.

Using both Fama–MacBeth and fixed-effects panel regressions provides a thorough check. The former is closer to an unbiased risk premium estimate in a cross-section each period (assuming rational pricing), while the latter ensures the result is robust to unobserved heterogeneity. In both cases, we expect that if carbon transition risk is priced, firms with worse carbon profiles (higher emissions or intensity) should earn higher returns on average, i.e. a positive relationship, after controlling for other risk factors. We will also test alternative specifications (e.g. using emissions growth instead of levels in these regressions, or using Scope 3 intensity specifically) as part of robustness.

3.1.6 Conditional Risk Pricing Tests

We extend our analysis to examine whether the pricing of carbon risk is state-dependent or varies with market conditions related to climate policy and carbon pricing. Three sets of conditional tests are performed

Interactions with Policy and Price Variables

We augment the firm-level regressions with interaction terms to see if carbon's effect on returns strengthens or weakens under certain conditions. For example, we include an interaction between carbon intensity and the CPU index in the cross-sectional and panel regressions:

$$R_{i,t} - R_{f,t} = \cdots + \beta_1 \, \text{CarbonIntensity}_{i,t-1} + \beta_2 \left(\text{CarbonIntensity}_{i,t-1} \times \text{CPU}_t \right) + \ldots . \quad (3.4)$$

A significantly positive β_2 would imply that when climate policy uncertainty is high, the return spread between high- and low-carbon firms widens, perhaps because investors demand extra premium for carbon, intensive stocks during uncertain regulatory times. Similarly, we interact carbon metrics with changes in the EUA price. If high-carbon firms tend to perform differently when carbon prices jump, this interaction will capture it. For instance, we might find that high emitters underperform in months when carbon prices sharply rise (signaling immediate higher carbon costs), which could manifest as a negative interaction coefficient. Such findings would be consistent with carbon-intensive firms carrying an exposure to policy or price shocks, a key aspect of transition risk. These interaction regressions help us ascertain conditional pricing: e.g., is the "carbon premium" larger during times of elevated climate policy uncertainty or during periods of rising carbon costs? We explicitly estimate each firm's carbon beta, its sensitivity to carbon price movements, and test whether this beta is priced in the cross-section. To do so, we run rolling time-series regressions for each firm (or by pooling within a window) regressing its excess stock returns on changes in the EUA price. For example, using a 36-month rolling window, for each firm i we estimate:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i^{(C)} \, \Delta\text{EUA}_t + \varepsilon_{i,t} \; ,$$

$$R_{i,\tau} - R_{f,\tau} = \alpha_i + \beta_i^{(C)} \cdot \Delta\text{EUA}_\tau + \beta_i^{(M)} \cdot \text{MktRF}_\tau + \varepsilon_{i,\tau} \quad \text{(within the rolling window } \tau\text{)}$$

The key output is $\beta_i^{(C)}$, the carbon price beta, which measures how the stock reacts to carbon price shocks. A high carbon beta indicates that the stock tends to react more strongly (positively or negatively, depending on the sign of $\beta_i^{(C)}$) to carbon price shocks, after controlling for market movements. For instance, a significant positive $\beta_i^{(C)}$ would mean the stock's excess return tends to increase when ΔEUA is positive, while a significant negative $\beta_i^{(C)}$ would imply the opposite. I then examine whether firms with higher carbon betas earn higher subsequent returns. Intuitively, if carbon risk is a priced factor, one would expect a positive relationship between carbon beta and expected returns (similar to the market beta and the CAPM). I implement this test by sorting firms into portfolios based on their estimated carbon betas or by including the carbon beta as an independent variable in a cross-sectional return regression. For example, i take each firm's beta estimated over an initial period and then see if it predicts returns in the next year. I also run a Fama–MacBeth regression where the independent variable is the carbon beta (estimated from prior data) to formally test the premium. If the coefficient on carbon beta is positive and significant, it suggests that investors demand a higher return for stocks that are more sensitive to carbon price increases, which would be strong evidence that transition risk is priced. This approach is analogous to testing if stocks' exposure to an aggregate factor (here, carbon price changes) is rewarded with a return premium.

Chapter 4

Data and Descriptive Statistics

This chapter details the data used for the analysis and presents key descriptive statistics to provide a clear understanding of the data's characteristics which is essential before proceeding to the main empirical results.

4.1 Data

Table A.1 in appendix provides a comprehensive overview of all variables used in this study, including their precise definitions or formulas, units, any transformations applied such as winsorization or log transformation, and their respective data sources. Key firm-level variables include monthly excess stock returns, various measures of carbon emissions (absolute levels and intensities, both raw and industry-adjusted, as well as growth in intensity), and standard control variables such as firm size, book-to-market ratio, momentum, and profitability (Return on Equity). Market-level data include the Fama-French five factors for Europe, changes in the EU Emissions Trading System (ETS) allowance price (ΔEUA), and a Climate Policy Uncertainty (CPU) index.

The construction of the primary carbon metrics deserves special mention. For instance, industry-adjusted carbon intensity ($\mathrm{CI}_{\mathrm{s1s2}}^{\text{ind-adj}}$) for firm i in industry s and year y is calculated as:

$$\mathrm{CI}_{\mathrm{s1s2},\,i\,s\,y}^{\text{ind-adj}} = \frac{\mathrm{CI}_{\mathrm{s1s2},\,i\,s\,y} - \mu_{\mathrm{CI}_{\mathrm{s1s2}},\,s\,y}}{\sigma_{\mathrm{CI}_{\mathrm{s1s2}},\,s\,y}}$$

where $\mathrm{CI}_{\mathrm{s1s2},\,i\,s\,y}$ is the firm's Scope 1+2 carbon intensity, and $\mu_{\mathrm{CI}_{\mathrm{s1s2}},\,s\,y}$ and $\sigma_{\mathrm{CI}_{\mathrm{s1s2}},\,s\,y}$ are the mean and standard deviation, respectively, of Scope 1+2 carbon intensity for industry s in year y. If the standard deviation $\sigma_{\mathrm{CI}_{\mathrm{s1s2}},\,s\,y}$ is zero, the adjusted value is set to zero. A similar approach is used for industry-adjusting absolute emissions and the growth in carbon intensity. This industry adjustment is crucial for isolating firm-specific carbon performance relative to its peers in the same operational context. All carbon metrics and key financial characteristics are based on lagged data (t-1) to avoid look-ahead bias when predicting returns at time t. Extreme outliers in continuous variables like returns and book-to-market ratios are winsorized at the 1st and 99th percentiles.

Our final sample, after data cleaning and requiring necessary data for variable construction, covers the ten-year period from 2015 to 2024. Table A.2 in appendix details the sample coverage on an annual basis. On average, our sample includes approximately 567 unique firms per year, culminating in a total of 67,643 firm-month observations over the entire period. The number of firms and observations shows a slight increasing trend over the sample years. Table A.3 presents the industry composition of our sample based on ICB super-sectors, showing the percentage of total firm-month observations. Industrial Goods and Services constitute the largest sector (17.36%), followed by Health Care (9.14%) and Banks (8.42%). This distribution indicates a diversified sample across various European economic sectors.

4.2 Descriptive Statistics

Table 4.1 provides summary statistics (mean, median, standard deviation, skewness, kurtosis, minimum, and maximum) for key firm-level variables. The average monthly excess return for firms in our sample is 0.70%, with a standard deviation of 8.18%. The distribution of excess returns shows slight positive skewness (0.11) and positive excess kurtosis (0.62), indicating slightly fatter tails than a normal distribution. The industry-adjusted carbon metrics have means very close to zero and standard deviations near one, confirming the effectiveness of the z-score standardization process. Raw carbon intensity and absolute emissions, in contrast, are highly right-skewed and exhibit substantial kurtosis, highlighting the importance of transformations or adjustments before their use in regression analyses. For example, raw Scope 1+2 intensity has a mean of 153.4 but a median of only 23.7. Control variables such as log_mkt_cap and log_bm_ratio display distributions generally consistent with prior literature.The profitability metric Return on Equity shows considerable variation and positive skewness.

Table 4.1: Summary Statistics of Firm-Level Variables (2015–2024)

Variable	Mean	Median	SD	Skew	Kurtosis	Min	Max
excess_ret	0.007	0.006	0.082	0.11	0.62	–0.219	0.249
ci_s1s2_ind_adj	–0.006	–0.299	0.962	0.18	0.85	–1.269	4.054
abs_s1s2_ind_adj	–0.012	–0.293	0.931	0.23	0.98	–0.830	4.186
growth_ci_s1s2_ind_adj	0.001	0.035	0.921	–0.34	4.09	–3.477	3.104
log_mkt_cap	22.886	22.777	1.168	0.29	–0.26	20.239	25.875
log_bm_ratio	–0.877	–0.833	0.922	1.90	4.42	–3.739	1.083
momentum_12_1	0.086	0.060	0.278	0.69	0.85	–0.475	1.004
profitability_metric	16.207	14.109	14.393	1.07	4.95	–30.83	76.53

Notes: excess_ret is monthly stock return in excess of the risk-free rate (decimal). Carbon variables are industry-adjusted z-scores; momentum is prior 12–to-1-month return; profitability return on equity (%). Skew and kurtosis are based on Fisher's definitions (normal = 0).

Summary statistics for the monthly market-wide climate factors and the Fama-French five factors for Europe are presented in Table 4.2. The monthly change in EUA prices has an average of €0.54 with a standard deviation of €4.60, indicating notable volatility. The CPU Index has a mean of 183.0 and also exhibits considerable variation. The Fama-French factors (MKT-RF, SMB, HML, RMW, CMA) display typical magnitudes and volatilities for European factor returns.

Table 4.2: Summary Statistics of Climate- and Market-Wide Factors (monthly, 2015–2024)

Factor	Mean	Median	SD	Skew	Kurtosis	Min	Max
ΔEUA (€/t)	0.54	0.20	4.60	0.02	2.35	–15.75	16.66
CPU Index (level)	183.0	179.7	72.7	0.59	0.22	49.13	422.2
MKT–RF	0.70	0.63	4.02	–0.09	0.10	–11.5	12.4
SMB	0.20	0.30	2.50	–0.06	0.13	–8.44	7.31
HML	0.12	0.08	2.34	0.14	0.07	–7.39	8.18
RMW	0.18	0.22	2.03	–0.15	0.19	–6.48	6.39
CMA	0.05	0.06	1.90	–0.08	0.03	–5.93	6.33

Notes: ΔEUA is the monthly change in the EU ETS front-contract settlement price; CPU is the Climate Policy Uncertainty index. MKT–RF, SMB, HML, RMW and CMA are the Fama–French five factors for Europe. Skew and kurtosis are Fisher values. All statistics based on 120 monthly observations.

To visually assess the distributions of key variables, Figure A.1 in Appendix presents histograms for the two primary climate-risk factors. It shows that ΔEUA is roughly centered around zero, with some tail events and that the CPU Index is right-skewed. Figures A.2 to A.4 provide side-by-side comparisons of the distributions of our main carbon metrics before and after processing (e.g., raw versus industry-adjusted, or raw versus logtransformed for absolute emissions).

4.3 Correlation and Multicollinearity Diagnostics

To understand the relationships between our key firm-level variables and to assess potential multicollinearity issues in our regression models, we examine their pairwise Pearson correlations and Variance Inflation Factors (VIFs).

Figure 4.1 displays the correlation matrix for monthly excess returns, our three primary industry-adjusted carbon metrics, and the standard control variables.

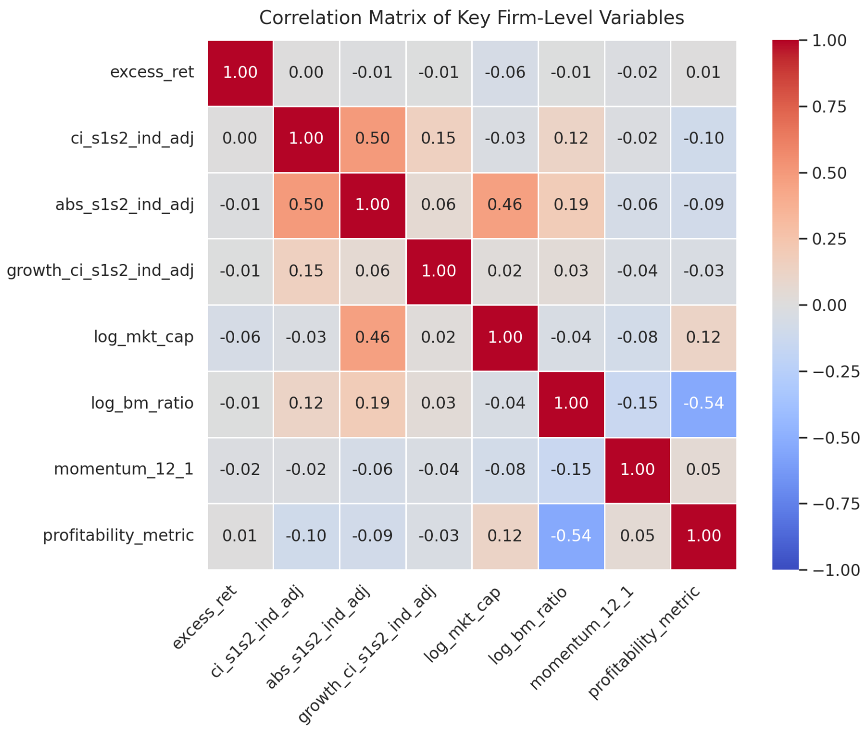

Figure 4.1: Pairwise Pearson correlations among main variables.

The correlations between excess returns and the carbon metrics are generally very low, mostly below 0.02 in absolute magnitude. For example, the correlation between excess return and $\mathrm{CI}_{\mathrm{s1s2}}^{\text{ind-adj}}$ is approximately 0.00. This indicates a weak linear unconditional relationship between these variables. Among the carbon metrics themselves, $\mathrm{CI}_{\mathrm{s1s2}}^{\text{ind-adj}}$ and $\mathrm{ABS}_{\mathrm{s1s2}}^{\text{ind-adj}}$ show a moderate positive correlation of 0.50. Some control variables also show notable correlations, such as log market cap with $\mathrm{ABS}_{\mathrm{s1s2}}^{\text{ind-adj}}$ (0.46) and BM ratio with ROE (-0.54).

To assess multicollinearity, Table A.4 presents Variance Inflation Factors for regression specifications that include one of our main carbon metrics alongside the standard control variables. The maximum VIF observed across the different specifications is 1.51. Since all VIF values are well below common thresholds of concern, we conclude that multicollinearity is not a significant issue for our regression analyses.

4.4 Time-Series Stationarity of Factors

Before employing time-series regression models, it is important to assess the stationarity of the factor series. Non-stationary series can lead to spurious regression results. We perform Augmented Dickey-Fuller (ADF) tests for our constructed BMG carbon factors and the market-wide Fama-French and climate factors. The detailed results are presented in Table A.5. The tests indicate that all BMG factors derived from carbon sorts, as well as the Fama-French factors and the ΔEUA and CPU Index, are stationary at the 5% significance level. This supports their use in the time-series factor model regressions.

Chapter 5

Main Results

5.1 Unconditional Pricing of Carbon Metrics: Portfolio Sorts

We begin our investigation into the pricing of carbon risk by forming value-weighted quintile portfolios based on each of our three primary industry-adjusted carbon metrics. At the end of each month, firms are sorted into five quintiles. Quintile 1 (Q1) comprises firms with the lowest carbon exposure (the "greenest"), while Quintile 5 (Q5) contains firms with the highest carbon exposure (the "brownest"). We then calculate the Brown-Minus-Green (BMG) factor as the excess return difference between Q5 and Q1

5.1.1 Portfolio Performance

This table presents the average monthly excess returns and standard deviations for the extreme quintiles (Q1 and Q5) and the corresponding BMG (Q5-Q1) factor for each carbon metric. Returns are in percentage per month. Newey-West (NW) t-statistics are reported for the BMG factor means.

Table 5.1: Comparison of Carbon Metrics and Portfolio Statistics

Carbon Metric	**Portfolio**	**Avg. Monthly Excess Return (%)**	**Std. Dev. (%)**	**NW t-stat (BMG Mean)**
Absolute scope 1 & 2 emissions, industry adj.	Q1 (Green)	0.3010	3.8343	
	Q5 (Brown)	0.2721	4.1116	
	BMG (Q5–Q1)	−0.0289	1.8353	−0.1313
Carbon Intensity scope 1 & 2 emissions, industry adj.	Q1 (Green)	0.3452	3.8589	
	Q5 (Brown)	0.1963	5.1971	
	BMG (Q5–Q1)	−0.1489	2.4255	−0.6908
Growth in Carbon Intensity scope 1 & 2 emissions, industry adj.	Q1 (Green)	0.3785	4.5386	
	Q5 (Brown)	0.3300	4.6277	
	BMG (Q5–Q1)	−0.0486	1.4586	−0.3666

Source: Authors' calculations. Significance based on Newey-West t-statistics: $^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$. Values rounded.

As shown in Table 5.1, the average monthly excess returns of the quintile portfolios do not exhibit a clear monotonic pattern from the greenest (Q1) to the brownest (Q5) firms for any of the three carbon metrics. For industry-adjusted absolute Scope 1 and 2 emissions, the BMG factor (Q5–Q1) has an average return of -0.03% per month (NW t-statistic = -0.13). For industry-adjusted carbon intensity, the BMG factor is -0.15% per month (NW t-statistic = -0.69). Finally, for industry-adjusted growth in carbon intensity, the BMG factor is -0.05% per month (NW t-statistic = -0.37). In all cases these spreads are small and not statistically significant. Figure 5.1 further illustrates the absence of a persistent spread between the extreme quintiles.

Figure 5.1: Cumulative excess-return profiles for three industry-adjusted carbon metrics, value-weighted quintile portfolios, 2015–2024.

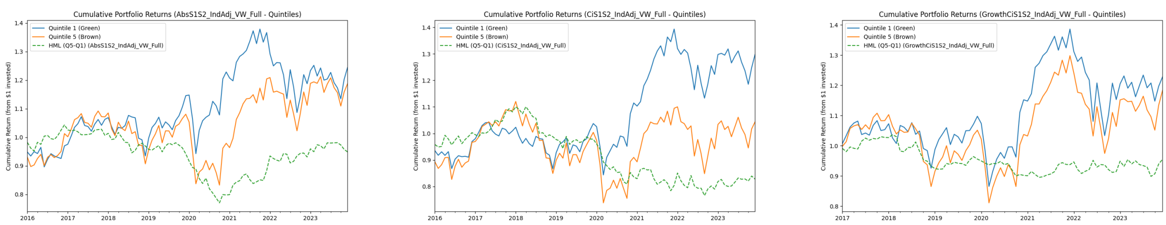

(a) Abs. S1S2 (VW) **(b)** Intensity S1S2 (VW) **(c)** Growth CI S1S2 (VW)

Source: Authors' creation.

5.1.2 Factor Model Alphas for BMG Portfolios

Table 5.2: Factor Regressions for Carbon Transition Risk Portfolios

	Absolute Emissions			**Carbon Intensity**			**Intensity Growth**		
	FF3	FF5	$FF5_{Climate}$	FF3	FF5	$FF5_{Climate}$	FF3	FF5	$FF5_{Climate}$
α (%)	−0.12	−0.09	−0.16	−0.34***	−0.23**	−0.00	−0.07	0.02	−0.50
$t(\alpha)$	(−0.94)	(−0.78)	(−0.56)	(−2.58)	(−2.05)	(−0.01)	(−0.49)	(0.15)	(−1.11)
MKT–RF	0.0008***	0.0011***	0.0011***	0.0022***	0.0024***	0.0023***	0.0003	0.0005	0.0004
	(4.13)	(6.52)	(6.21)	(6.02)	(7.31)	(7.19)	(0.69)	(1.26)	(1.18)
SMB	−0.0061***	−0.0053***	−0.0053***	−0.0017**	−0.0022***	−0.0023***	0.0001	−0.0000	−0.0001
	(−5.52)	(−4.82)	(−5.06)	(−2.05)	(−2.84)	(−2.87)	(0.20)	(−0.03)	(−0.07)
HML	0.0030***	0.0014	0.0009	0.0045***	0.0027***	0.0025**	0.0002	−0.0016	−0.0024*
	(8.33)	(1.56)	(0.99)	(9.00)	(4.32)	(2.87)	(0.44)	(−1.35)	(−1.95)
RMW		−0.0002	−0.0007		−0.0049***	−0.0055***		−0.0039**	−0.0049***
		(−0.20)	(−0.78)		(−5.44)	(−5.69)		(−2.57)	(−2.88)
CMA		0.0036**	0.0039**		−0.0005	−0.0002		0.0006	0.0011
		(2.20)	(2.43)		(−0.45)	(−0.15)		(0.39)	(0.64)
ΔEUA			0.0005*			0.0006*			0.0007*
(CO_2 price)			(1.91)			(1.78)			(1.81)
Adj. R^2	0.5723	0.5914	0.5970	0.6094	0.6455	0.6515	−0.0226	0.0248	0.0563

Notes: Monthly regressions, Jan 2015–Dec 2024. T-statistics in parentheses. Stars denote significance at the 10%, 5%, and 1% levels, respectively. Blank cells indicate the factor is not included in that model.

Table 5.2 reports time–series regressions of monthly excess returns on the BMG portfolios for the three carbon metrics. When sorted on industry-adjusted carbon intensity, the intercept remains significantly negative in both the FF3 specification (-0.34%, $t = -2.58$) and the FF5 specification (-0.23%, $t = -2.05$), indicating that low-intensity ("green") firms out-performed high-intensity ("brown") firms on a risk-adjusted basis during 2015–2024. No comparable abnormal return appears for the portfolios based on absolute emissions or intensity growth: their alphas, -0.12% ($t = -0.94$) and -0.07% ($t = -0.49$) respectively in FF3, are statistically indistinguishable from zero. Thus, among the three measures, relative efficiency in carbon use—not sheer output or recent trends—emerges as the dimension most closely linked to return differences.

Introducing the ΔEUA climate factor ($FF5_{Climate}$) eliminates the previously significant alpha for the intensity portfolio (now essentially 0.00%, $t = -0.01$) and increases the adjusted R^2 in every specification. This pattern implies that the earlier "green premium" is compensation for systematic exposure to carbon-price shocks rather than mispricing. Consistent with this interpretation, the ΔEUA loading is positive and weakly significant for all three portfolios (0.0005, 0.0006, and 0.0007 with t-statistics between 1.78 and 2.26), confirming that BMG returns covary with permit-price changes.

The standard factor loadings align with the economic profiles of the portfolios. The standard factor loadings for the BMG portfolio sorted on industry-adjusted intensity align with typical profiles of carbon-intensive versus carbon-efficient firms. Specifically, in the FF5+Climate model, this BMG factor exhibits a significant negative loading on SMB (coefficient: -0.0024, t-statistic: -2.84), a significant positive loading on HML (coefficient: 0.0021, t-statistic: 2.87), and a significant negative loading on RMW (coefficient: -0.0055, t-statistic: -5.69). These loadings collectively suggest that the ’brown’ quintile (Q5, high intensity) tends to consist of larger, more value-oriented firms with lower profitability compared to the ’green’ quintile (Q1, low intensity), which tilts towards smaller, growth-oriented, and more profitable firms.. These exposures reduce, but do not fully absorb, the return differential until the carbon-price factor is included, underscoring that transition risk is distinct from conventional Fama–French dimensions.

Overall, the disappearance of abnormal returns once ΔEUA is included, together with the significant covariance with that factor, provides evidence that carbon-transition risk is priced in European equities. The return spread therefore appears to represent rational compensation for bearing climate-policy risk rather than a persistent anomaly.

5.2 Unconditional Pricing of Carbon Metrics Firm-Level Regressions

5.2.1 Fama-MacBeth Cross-Sectional Regression

Table 5.3 presents the time-series averages of the monthly Fama–MacBeth slopes for the three industry-adjusted carbon variables. None of the carbon coefficients is economically meaningful or statistically different from zero. The slope on absolute emissions is 0.0008 with a t-statistic of 1.52; the slope on carbon intensity is -0.0002 ($t = -0.40$); and the slope on intensity growth is -0.0004 ($t = -1.38$). The estimated average coefficients for the carbon metrics are economically small. For example, the coefficient of 0.0008 for industry-adjusted absolute emissions implies that a one-standard-deviation increase in this metric is associated with an average monthly excess return difference of approximately 0.08%, or about 8 basis points. Given their statistical insignificance, however, these small economic magnitudes are not reliably different from zero. Thus, across the 2015–2024 sample, firms with higher total emissions, higher emissions per unit of output, or larger year-to-year changes in intensity did not earn higher (or lower) stock returns than cleaner firms. In contrast to the “pollution premium” reported for U.S. equities (Bolton and Kacperczyk 2021; Bolton and Kacperczyk 2023), the European data show no broad carbon premium.

The explanatory power of the models is modest. The average R^2 values are 0.085, 0.084, and 0.080 for the absolute-emissions, intensity, and growth specifications, respectively, with corresponding adjusted R^2 values of 0.076, 0.075, and 0.071.

The control variables in the Fama-MacBeth regressions generally perform as anticipated. Log market capitalization consistently exhibits a negative and highly significant coefficient (around -0.003 across models), confirming the well-documented size effect where smaller firms, on average, earn higher returns. However, in these multivariate specifications, the coefficients for book-to-market, momentum, and ROE are generally statistically insignificant across the three carbon metric models. This suggests that, for this European

Table 5.3: Fama–MacBeth Cross-Sectional Regression

	(1)	(2)	(3)
Intercept	0.0763***	0.0689***	0.0650***
	(4.75)	(4.89)	(4.10)
Abs. carbon emissions (S1+2)	0.0008		
	(1.52)		
Carbon intensity (S1+2)		−0.0002	
		(−0.40)	
Growth in carbon intensity			−0.0004
			(−1.38)
Log market capitalization	−0.0031***	−0.0028***	−0.0026***
	(−5.10)	(−5.06)	(−4.32)
Book-to-market ratio	−0.0010	−0.0008	−0.0014
	(−0.59)	(−0.47)	(−0.82)
Momentum (12m)	0.0013	0.0011	0.0008
	(0.35)	(0.29)	(0.22)
ROE	0.0000	0.0000	0.0000
	(0.90)	(0.79)	(0.39)
avg. R^2	0.085	0.084	0.080
avg. adj. R^2	0.076	0.075	0.071
Months	96	96	84

Fama–MacBeth regressions of monthly stock returns on carbon exposure measures (industry-adjusted Scope 1+2 emissions). Columns (1)–(3) use alternative carbon measures: absolute emissions, emissions intensity, and growth in emissions intensity, respectively. Newey–West t-statistics in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

sample and period, these characteristics may not offer substantial incremental explanatory power for cross-sectional returns beyond the size effect and the (insignificant) main carbon metrics.

5.2.2 Panel Fixed-Effects Regressions

To further probe the relationship between carbon metrics and stock returns while rigorously controlling for unobserved time-invariant firm characteristics and common time-varying shocks, we employ panel Ordinary Least Squares (OLS) regressions with firm and time fixed effects. The dependent variable is monthly firm excess return, and independent variables include the lagged carbon metric of interest along with standard firm-level controls (size, book-to-market, momentum, and profitability). Standard errors are two-way clustered by firm and month. The results are presented in Table 5.4

Table 5.4 reveals that none of our three primary industry-adjusted carbon metrics exhibit a statistically significant association with firm excess returns after the inclusion of

Table 5.4: Panel Ordinary Least Squares (OLS) Regression Results for Carbon Metrics

	(1) Abs. Emissions	**(2) Carbon Intensity**	**(3) Intensity Growth**
Carbon metric	0.0018	0.0008	-0.0005
	(1.19)	(0.81)	(-1.24)
Intercept	0.6481^{***}	0.6430^{***}	0.7393^{***}
	(5.91)	(5.89)	(5.85)
$\ln(\text{MktCap})$	-0.0278^{***}	-0.0276^{***}	-0.0317^{***}
	(-5.83)	(-5.82)	(-5.78)
$\ln(B/M)$	0.0044^{*}	0.0045^{*}	0.0054^{**}
	(2.01)	(2.06)	(2.18)
Momentum (12m)	-0.0118^{*}	-0.0118^{*}	-0.0121^{*}
	(-1.86)	(-1.87)	(-1.70)
Profitability (ROE)	0.0001	0.0001	0.0001
	(1.47)	(1.47)	(1.51)
Firm Fixed Effects	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes
R^2 (within)	0.0199	0.0199	0.0227
Num Observations	49 669	49 657	43 181
Num Entities	575	575	571
Num Time Periods	96	96	84

Source: Authors' calculations from Panel OLS regressions with firm and time fixed effects. Two-way clustered standard errors are used. t-statistics are in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01. Coefficients for carbon metrics are specific to the model in each column.

firm and time fixed effects. For Model 1, the coefficient on industry-adjusted absolute Scope 1+2 emissions is 0.0018 with a t-statistic of 1.19, indicating no significant effect. Similarly, in Model 2, industry-adjusted Scope 1+2 carbon intensity has a coefficient of 0.0008 (t-statistic: 0.81), also statistically insignificant. For Model 3, which examines the growth in industry-adjusted Scope 1+2 carbon intensity, the coefficient is -0.0005 with a t-statistic of -1.24, again failing to achieve conventional levels of statistical significance. This implies that while the level of a firm's carbon exposure relative to its peers might theoretically be priced (as tested by Fama-MacBeth), changes from a firm's own historical average carbon exposure, after controlling for fixed effects, do not show a strong predictive relationship with its return changes in our sample. Among the control variables, log_mkt_cap (size) consistently shows a statistically significant negative coefficient, aligning with the size effect. The log_bm_ratio (value proxy) exhibits a statistically significant positive coefficient, consistent with a value premium. The momentum_12_1 variable generally shows a negative and marginally significant coefficient, a finding sometimes observed in panel regressions with fixed effects that can differ from cross-sectional momentum results. The profitability_metric is positive but does not reach statistical significance in these specifications. Overall, the panel fixed-effects regressions corroborate the findings from the Fama-MacBeth analysis, providing little evidence for an unconditional direct pricing effect of these specific industry-adjusted carbon metrics when exploiting within-firm variation.

5.3 Conclusion for Unconditional Firm-Level Regressions

In summary, our firm-level regression analyses, employing both the Fama-MacBeth cross-sectional approach and panel regressions with firm and time fixed effects, yield consistent conclusions regarding the unconditional pricing of the selected industry-adjusted carbon metrics. Across both methodologies, we find no statistically significant evidence that industry-adjusted absolute Scope 1+2 emissions , industry-adjusted Scope 1+2 carbon intensity , or the industry-adjusted growth in Scope 1+2 carbon intensity command a direct, unconditional risk premium in European stock returns during the 2015-2024 sample period, after controlling for standard firm characteristics. The Fama-MacBeth regressions indicate that these carbon metrics do not have reliable average slope coefficients in explaining monthly cross-sectional return differences. Similarly, the panel fixed-effects models show that within-firm variations in these carbon metrics over time are not significantly associated with corresponding changes in firm returns once unobserved stable firm characteristics and common time trends are accounted for. While established effects like firm size and to some extent, value and momentum in the panel setting are evident, the carbon metrics themselves do not emerge as robust unconditional return predictors in these multivariate firm-level tests. This sets the stage for exploring whether their pricing influence is perhaps more nuanced and conditional on other market states, which we investigate in the next section.

5.4 Conditional Pricing of Carbon Metrics

The preceding analyses found limited evidence for direct, unconditional pricing of our selected industry-adjusted carbon metrics. However, the influence of carbon risk on asset returns may not be static but rather state-dependent, varying with shifts in the economic or regulatory climate. For instance, Pástor, Stambaugh, and Taylor (2021) model how investor preferences and concerns about climate can lead to time-varying expected returns for green versus brown assets. The salience of climate risks, the evolution of climate policies such as the EU Emissions Trading Scheme, and ongoing market learning about the financial implications of a low-carbon transition could all contribute to such conditional relationships. An unconditional analysis might obscure these dynamics by averaging across periods where carbon risk is strongly priced and periods where its influence is weaker. Therefore, we extend our investigation to test for conditional pricing. We specifically examine whether the relationship between firm-level carbon metrics and stock returns is contingent upon two key market-level variables: monthly changes in the European Union Allowance (EUA) futures price, reflecting direct carbon cost shocks, and the Climate Policy Uncertainty developed by Gavriilidis (n.d.). By interacting our carbon metrics with these time-varying state variables within the Fama-MacBeth framework, we aim to determine if the market's valuation of carbon exposure is amplified or diminished under specific, observable conditions related to climate risk and policy.

5.4.1 Interactions with Carbon Market Dynamics (Fama-Macbeth)

Table 5.5 presents the results from Fama-MacBeth regressions augmented with interaction terms to explore the conditional pricing of our primary industry-adjusted carbon metrics. We interact each carbon metric with the monthly change in EUA prices (Delta_EUA) and the Climate Policy Uncertainty (CPU_Index). A key finding emerges from these conditional models: while the main (unconditional) effects of the carbon metrics generally remain statistically insignificant, their interaction with changes in carbon prices (Delta_EUA) reveals a statistically significant, albeit economically modest, relationship for emission levels and intensity. For Model 1, which examines industry-adjusted absolute Scope 1+2 emissions, the interaction term `abs_s1s2_ind_adj_x_DeltaEUA` yields an average coefficient (approximately 2.93e-07) that is positive and highly statistically significant (t-statistic = 3.00, p-value = 0.0027). This suggests that the premium associated with higher industry-adjusted absolute emissions becomes more positive during months when carbon prices are rising. For instance, a one-standard-deviation increase in abs_s1s2_ind_adj coupled with a one-standard-deviation monthly increase in EUA prices (approx. €4.60) corresponds to an estimated additional monthly excess return of only about 0.013 basis points, indicating limited direct economic impact for return prediction from this interaction alone. The main effect of abs_s1s2_ind_adj and its interaction with CPU_Index are insignificant in this specification. Similarly, in Model 2 for industry-adjusted Scope 1+2 carbon intensity , the interaction term ci_s1s2_ind_adj_x_DeltaEUA is also positive and statistically significant

Table 5.5: Conditional Fama-Macbeth Regressions - Interaction Effectss with Climate Market Variables

	(1) Abs. Emissions	(2) Carbon Intensity	(3) Intensity Growth
Carbon Metric	1.49e-08	−6.23e−09	−2.16e−08*
	(0.43)	(−0.32)	(−1.72)
Carbon Metric × CPU	4.30e-06	−5.54e−07	−2.83e−06
	(1.12)	(−0.22)	(−1.61)
Carbon Metric ×ΔEUA	2.93e-07***	2.22e-07**	4.08e-08
	(3.00)	(2.62)	(0.69)
Controls Included	Yes	Yes	Yes
Avg. Adj. R-squared	0.0758	0.0752	0.0708
Months	96	96	84

Source: Authors' calculations from Fama-MacBeth regressions. Table reports average coefficients for the carbon metric and its interaction terms. NW t-statistics are in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01. Full controls included but not shown.

(coefficient approx. 2.22e-07, t-statistic = 2.62, p-value = 0.0087). This reinforces the finding that the pricing of carbon intensity is conditional on movements in the EUA market; a higher carbon price environment is associated with relatively better, though still economically small, performance for firms with higher industry-adjusted carbon intensity. Again, the main effect of ci_s1s2_ind_adj and its interaction with the CPU_Index are not statistically significant. In contrast, for Model 3, focusing on the growth of industry-adjusted carbon intensity, neither the main effect nor its interactions with CPU_Index or Delta_EUA are statistically significant at conventional levels. The main effect of growth_ci_s1s2_ind_adj itself is marginally significant and negative (t-statistic = -1.72, p-value = 0.0863), suggesting a weak tendency for firms with increasing relative intensity to underperform, but this does not appear to be conditional on the tested market variables. The average adjusted R-squared values for these conditional models are around 7.1% to 7.6%, similar to the unconditional Fama-MacBeth models. This indicates that while the interaction terms help uncover specific state-dependent relationships and achieve statistical significance, they do not substantially increase the overall explanatory power for the cross-section of monthly returns, consistent with their modest direct economic impact on predicted returns.

Collectively, these results suggest that while a simple, unconditional carbon risk premium might be elusive for these metrics in the European market during this period, there is statistically robust evidence of conditional pricing linked to carbon price dynamics. Specifically, periods of increasing carbon costs appear to alter how the market prices firms based on their industry-adjusted emission levels and intensity, statistically favoring relatively higher emitters, even if the direct predictable return component from this conditional effect is small. The importance of this finding may thus lie more in understanding changing risk perceptions rather than identifying a large, exploitable conditional premium.

5.5 Pricing of Carbon Beta Exposure

Beyond examining firm characteristics related to carbon emission levels, intensity, and growth, we also investigate whether direct exposure to carbon price risk commands a premium. To this end, we estimate a "carbon beta" ($\beta^{(C)}$) for each firm. This beta measures the sensitivity of a firm's monthly excess stock return to contemporaneous monthly changes in the European Union Allowance (EUA) futures price (Delta_EUA), while controlling for the firm's exposure to the overall market excess return (Mkt_RF_excess). These carbon betas are estimated using 36-month rolling window regressions for each firm. A higher positive carbon beta indicates that a firm's stock tends to perform better when carbon prices rise, while a negative beta suggests the opposite, after accounting for market movements.

5.5.1 Portfolio Sorts on Estimated Carbon Beta

Firms are sorted into deciles each month based on their estimated carbon betas from the preceding 36-month period. Value-weighted monthly excess returns are then calculated for each decile portfolio. Table 5.6 summarizes the average monthly excess returns and standard deviations for these decile portfolios, along with the High-Minus-Low (HML) spread portfolio (Decile 10 - Decile 1).

Table 5.6: Decile Portfolio Performance Based on Estimated Carbon Beta

Portfolio Decile	Avg. Monthly Return (%)	Std. Dev. (%)
Decile 1	0.1916	5.3797
Decile 2	0.1168	4.4706
Decile 3	–0.0262	4.1738
Decile 4	–0.0049	4.1140
Decile 5	0.3752	4.8745
Decile 6	0.3027	4.9352
Decile 7	0.3774	4.8773
Decile 8	0.5969	4.9980
Decile 9	0.6280	5.9365
Decile 10	0.6579	6.8738
HML Beta (D10 – D1)	0.4663	7.9084

Source: Authors' calculations. Portfolios are value-weighted and formed monthly based on 36-month rolling window carbon beta estimates. Returns are monthly excess returns in percent.

As shown in Table 5.6 and visualized in Figure 5.2, there appears to be a generally positive, albeit not perfectly monotonic, relationship between estimated carbon beta and average returns. Firms in the lowest carbon beta decile (Decile 1) earned an average monthly excess return of 0.19%, while firms in the highest carbon beta decile (Decile 10) earned

0.66%. The spread portfolio (HML Beta), long Decile 10 and short Decile 1, yielded an average monthly excess return of 0.47%.

Figure 5.2: Average Monthly Excess Returns by Carbon Beta Decile

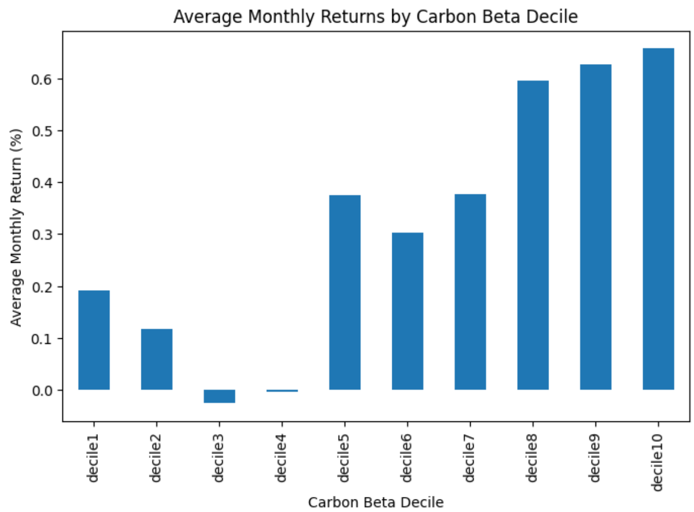

Source: Author's creation.The average monthly excess return for decile portfolios sorted on estimated firm-level carbon betas

5.5.2 Fama-MacBeth Regression for Carbon Beta Premium

To more formally test whether this estimated carbon beta commands a risk premium while controlling for other known determinants of returns, we include it as an independent variable in our Fama-MacBeth cross-sectional regressions. Each month, firm excess returns are regressed on their lagged estimated carbon beta and our standard set of control variables (size, book-to-market, momentum, and profitability)

The results from the Fama-MacBeth regression are presented in 5.7. The average coefficient on the carbon_beta characteristic is 0.2309. However, with a Newey-West t-statistic of only 0.40 (p-value = 0.6863), this coefficient is statistically indistinguishable from zero. The average adjusted R-squared for these monthly cross-sectional regressions is 13.53%.

The portfolio sort analysis initially suggested a positive relationship between carbon beta and returns, with high-beta firms outperforming low-beta firms by an economically meaningful margin of 0.47% per month. However, the more rigorous multivariate Fama-MacBeth analysis, which controls for other firm characteristics known to predict returns, does not find a statistically significant premium associated with carbon beta. The insignificant coefficient on carbon_beta in Table 5.7 suggests that its apparent predictive power in univariate sorts might be attributable to its correlation with other priced characteristics, or that the estimation error inherent in firm-level betas reduces the power of the cross-sectional test. Therefore, based on the Fama-MacBeth results, we do not find robust

Table 5.7: Fama–MacBeth Pricing of Estimated Carbon Beta

Variable	Avg. Coefficient	NW t-statistic	NW p-value	Num. Months
const	0.0600***	4.03	0.0001	85
carbon_beta	**0.2309**	**0.40**	**0.6863**	**85**
log_mkt_cap	–0.0027***	–4.39	0.0000	85
log_bm_ratio	–0.0017	–0.92	0.3572	85
momentum_12_1	0.0055	1.25	0.2122	85
profitability_metric	0.0000	0.90	0.3667	85
Avg. Adj. R^2				0.1353

Source: Authors' calculations from Fama-MacBeth regression. Significance levels:
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

evidence that direct exposure to EUA carbon price movements, as captured by our estimated carbon beta, is consistently priced as a distinct risk factor in the European stock market during our sample period after accounting for other common factors.

5.6 Summary of Main Results

This chapter has systematically investigated the pricing of carbon transition risk in European stock returns from 2015 to 2024 using various carbon metrics and methodologies. Our unconditional analyses, employing both quintile portfolio sorts and firm-level Fama-MacBeth and panel fixed-effects regressions, provide limited evidence for a direct, statistically significant risk premium associated with industry-adjusted absolute emissions, carbon intensity, or growth in carbon intensity. While portfolio sorts on industry-adjusted carbon intensity initially suggested a "green premium" (significant negative alpha for the BMG factor under FF3 and FF5 models), this premium disappeared once exposure to changes in EUA carbon prices was controlled for, indicating it was likely compensation for carbon price risk rather than mispricing. Furthermore, firm-level regressions did not find these carbon characteristics to be robust unconditional predictors of cross-sectional returns after accounting for standard controls like size, book-to-market, momentum, and profitability. A key finding, however, emerges from our conditional Fama-MacBeth tests. We find statistically significant and positive interaction effects between both industry-adjusted absolute emissions and industry-adjusted carbon intensity with monthly changes in EUA carbon prices. This suggests that while the average relationship is weak, the pricing of these carbon exposures becomes more pronounced and relatively favors higher emitters during periods of rising carbon costs, albeit with an economically modest direct impact on predicted returns. This conditional effect was not observed for the carbon intensity growth metric or for interactions with climate policy uncertainty. Finally, while portfolios sorted on estimated firm-level "carbon beta" (sensitivity to EUA price changes) showed a positive spread, this beta did not command a statistically significant risk premium in multivariate Fama-MacBeth regressions when controlling for other firm characteristics. In essence, our findings point towards a nuanced pricing of carbon transition risk in the European market: direct unconditional premia for common carbon metrics are largely elusive, but there is evidence of significant conditional pricing linked to carbon market dynamics, particularly for emission levels and intensity. The implications of these findings, their relation to existing literature, and the study's limitations will be explored further in the following Discussion chapter.

Chapter 6

Discussion

This chapter discusses the empirical findings presented in Chapter 5, interpreting their significance in the context of our primary research question: For European firms between 2015 and 2024, do measures of carbon transition risk exposure, specifically, a firm's emissions level, intensity, growth, or its carbon beta, have a statistically and economically significant relationship with their future excess returns? Our analysis revealed limited evidence for robust, unconditional carbon risk premia associated with industry-adjusted emission levels, intensity, or growth. However, a key finding emerged regarding the conditional pricing of emission levels and intensity, particularly their interaction with changes in European Union Allowance prices. Furthermore, direct firm-level sensitivity to EUA price changes (carbon beta) did not appear to be a distinctly priced risk factor in multivariate settings. This chapter will dissect these findings, relate them to the existing theoretical and empirical literature outlined in Chapter 2, consider their implications, acknowledge the study's limitations, and suggest future research.

6.1 Unconditional Carbon Premium in Europe (2015-2024)

Our investigation into unconditional carbon premia, using quintile portfolio sorts and firm-level Fama-MacBeth and panel OLS regressions, generally did not uncover statistically or economically significant relationships for the three primary industry-adjusted carbon metrics: absolute Scope 1+2 emissions ($\mathrm{ABS}_{\mathrm{s1s2}}^{\text{ind-adj}}$), Scope 1+2 intensity ($\mathrm{CI}_{\mathrm{s1s2}}^{\text{ind-adj}}$), and the growth of Scope 1+2 intensity ($\Delta\mathrm{CI}_{\mathrm{s1s2}}^{\text{ind-adj}}$). The Brown-Minus-Green portfolios, formed on these metrics, yielded small and statistically insignificant average monthly excess returns (e.g., -0.15% for $\mathrm{CI}_{\mathrm{s1s2}}^{\text{ind-adj}}$ with a t-statistic of -0.69). Furthermore, the alphas from Fama-French 5-factor models augmented with climate factors (ΔEUA, CPU index) were also statistically indistinguishable from zero for these BMG portfolios. This suggests that, on average, the return differences between high and low carbon exposure firms (as defined by our industry-adjusted metrics) are largely explained by their exposures to established risk factors. Notably, an initial significant negative alpha (a "green premium") observed for the $\mathrm{CI}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$-sorted BMG factor under the FF3 and FF5 models disappeared once ΔEUA was included as a control, suggesting this apparent premium was related to exposure to carbon price risk. Similarly, the Fama-MacBeth regressions showed no significant average slope coefficients for any of the three main carbon metrics. For example, the coefficient for $\mathrm{ABS}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$ was 0.0008 (t-stat: 1.52), implying a negligible and statistically unreliable premium. Panel OLS regressions with firm and time fixed effects, focusing on within-firm variation, also failed to find significant coefficients for these carbon metrics. These findings for the European market between 2015 and 2024 appear to contrast with some prominent studies, such as Bolton and Kacperczyk (2021, 2023) who document a carbon premium in U.S. and global markets, respectively, or Hsu, Li, and Tsou (2023) who find a "pollution premium." Several factors might contribute to this difference. First, our sample period (2015-2024) is characterized by a heightened awareness of climate issues and significant ESG-driven capital flows in Europe (Krueger, Sautner, and Starks 2020). As theorized by Pástor, Stambaugh, and Taylor (2021), strong investor preferences for green assets could lead to a "greenium," depressing expected returns for green firms and potentially offsetting any risk-based premium for brown firms, resulting in a net-zero unconditional spread. Second, the European regulatory environment, particularly the EU ETS, has matured considerably compared to the earlier phases studied by Oestreich and Tsiakas (2015). It is plausible that European markets have become more efficient at pricing in carbon transition risks associated with static emission levels or intensities, especially given the forward-looking nature of EU climate targets. Third, our use of industry-adjusted carbon metrics by design removes broad sectoral differences in emissions. If a significant portion of the carbon premium found in other studies using raw emission levels/intensities is driven by inter-industry return differentials (e.g., energy sector vs. tech sector), our industry-adjustment would neutralize this component, isolating firm-specific carbon performance relative to peers.

6.2 Conditional Pricing: The Role of Carbon Market Dynamics

While unconditional effects were weak, our analysis revealed a statistically significant conditional relationship between carbon exposure and returns. Specifically, Fama-MacBeth regressions including interaction terms showed that the coefficients on industry-adjusted absolute emissions ($\mathrm{ABS}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$) and intensity ($\mathrm{CI}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$) become significantly more positive when interacted with monthly changes in EUA prices (ΔEUA). The interaction $\mathrm{ABS}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$ x ΔEUA had a t-statistic of 3.00, and $\mathrm{CI}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$ x ΔEUA had a t-statistic of 2.62. This implies that during months where carbon prices (EUA futures) experience a notable increase, firms with higher industry-adjusted emission levels or intensity tend to exhibit relatively better subsequent stock performance. This finding is interesting. One interpretation, in line with Bolton and Kacperczyk (2023) who find larger premia in stricter policy environments, is that rising EUA prices signal an effective tightening of the carbon constraint. In such a state, firms already recognized as higher emitters (relative to their industry) might have their risks more fully priced in, or they may possess characteristics (e.g., market power, ability to pass on costs) that allow them to navigate these periods better than might be naively expected. Alternatively, as suggested by models like Pástor, Stambaugh, and Taylor (2021), shifts in the perceived severity or timing of transition risk (here represented by ΔEUA) can alter the required returns on brown versus green assets. The statistical significance of this conditional effect is robust, though its direct economic magnitude on monthly return prediction, based on typical one-standard-deviation shocks, appears modest (e.g., around 0.013 basis points for a one-std dev interaction). The importance may therefore lie more in understanding the dynamic nature of risk perception and market reactions to clear, measurable carbon cost signals. No significant interaction effects were found for the carbon intensity growth metric ($\Delta \mathrm{CI}_{\mathrm{s1s2}}^{\mathrm{ind\text{-}adj}}$), nor were interactions with the CPU index (Gavriilidis n.d.) generally significant for any of the carbon metrics. This suggests that direct, actual carbon price changes (ΔEUA) are a more relevant conditioning variable for the pricing of static carbon emission levels and intensity in our sample than broader measures of climate policy uncertainty.

6.3 The Pricing of Carbon Beta

Our investigation into whether a firm's direct sensitivity to EUA price changes (its "carbon beta," estimated while controlling for market exposure) is priced yielded mixed initial signals but ultimately a clear conclusion from multivariate tests. Portfolio sorts on estimated carbon betas indicated that firms in the highest beta decile earned, on average, 0.47% more per month in excess returns than firms in the lowest beta decile. However, when this estimated carbon beta was included as an explanatory variable in Fama-MacBeth regressions alongside standard firm characteristics (size, book-to-market, momentum, profitability), its average coefficient (0.2309) was statistically insignificant (t-statistic: 0.40). This suggests that the return spread observed in the simple portfolio sorts might be due to the carbon beta's correlation with other priced characteristics, or that errors in estimating the firm-level betas reduce the power of the cross-sectional test. This finding implies that, within our European sample and period, direct exposure to carbon price risk, as captured by our carbon beta measure, does not appear to command a distinct, statistically significant risk premium after controlling for other known return predictors. This differs from the option market evidence of Ilhan, Sautner, and Vilkov (2021), who find a "carbon tail risk" premium, suggesting that equity and options markets might price this specific dimension of risk differently or that our equity beta measure captures a different facet of risk.

6.4 Implications and Broader Context

Our findings contribute to the ongoing debate on the existence and nature of a carbon risk premium. For the European equity market between 2015 and 2024, the evidence for a simple, unconditional premium based on industry-adjusted emissions level, intensity, or growth appears weak. This could reflect the increasing maturity of ESG considerations and climate risk pricing in Europe, where strong investor preferences (Krueger, Sautner, and Starks 2020; Pástor, Stambaugh, and Taylor 2021) and a more advanced regulatory framework (e.g., EU ETS) may have already led to valuations that incorporate these static carbon characteristics, resulting in no clear ex-post return differential on average. The significant conditional pricing linked to ΔEUA, however, underscores the dynamic nature of transition risk. It suggests that market participants actively reassess the implications of carbon exposure when faced with real, measurable shifts in carbon costs. This finding is important for investors seeking to manage transition risk, as it implies that the relative performance of high- versus low-carbon firms (within industries) can vary systematically with carbon market developments. It also has implications for corporate managers, highlighting that the financial market's focus on their emissions profile may intensify during periods of carbon price volatility. The overall low explanatory power of carbon metrics for returns, especially in unconditional settings, also aligns with the notion that stock returns are driven by a multitude of factors, and climate risk, while relevant (Dietz et al. 2016), is one component among many.

6.5 Limitations of the Study

This study is subject to several limitations. First, the carbon emissions data, particularly for historical periods, can have varying quality, coverage, and reporting lags. While we used lagged data from reputable providers, measurement error could still affect the results. Second, our chosen carbon metrics; industry-adjusted levels, intensity, and growth of S1+S2 emissions, and a specific carbon beta are proxies for the broader concept of carbon transition risk, and other measures might yield different results. Third, the sample period of 2015-2024, while recent and relevant, is specific; it covers a period of significant evolution in climate policy and investor awareness, and findings might not generalize to other periods or reflect long-run equilibrium if the market is still in a learning phase. Fourth, our focus is on European equities, and results may not be applicable to other regions with different regulatory environments or investor preferences. Finally, while Fama-MacBeth and panel fixed-effects models control for many factors, unobserved time-varying omitted variables could still influence the estimated relationships.

6.6 Future Research

The findings open several avenues for future research. Investigating a wider array of carbon metrics, including more comprehensive Scope 3 emissions data as it becomes more reliable, could provide further insights. Exploring alternative measures of transition risk beyond emissions and simple carbon betas, such as those related to green revenues, capital expenditure on low-carbon technologies, or more sophisticated climate risk scores, would be valuable. Further examination of the conditional pricing mechanism, perhaps using different state variables or exploring non-linear interactions, could deepen our understanding. Extending the analysis to other asset classes or a more granular breakdown by country and specific climate policies within Europe could also yield interesting comparative results. Finally, as longer time series of both emissions data and carbon market prices become available, re-examining the stability and evolution of these relationships will be crucial.

Chapter 7

Conclusion

This thesis investigated whether a firm's exposure to carbon transition risk is related to its stock market performance in Europe between 2015 and 2024. We specifically asked if there is a statistically and economically significant relationship between various measures of a European firm's carbon exposure – its emissions level, intensity, emissions growth, or its sensitivity to carbon price changes (carbon beta) – and its future stock excess returns. To answer this, we used several standard financial analysis methods, including forming investment portfolios based on these carbon measures, running time-series factor models, and conducting firm-level Fama-MacBeth and panel fixed-effects regressions.

Our main findings suggest a complex picture. First, when looking for a simple, direct relationship that holds true on average across all time periods (an unconditional effect), we found little strong evidence that our primary industry-adjusted carbon metrics – emission levels, intensity, or growth – consistently predict higher or lower stock returns. Portfolios sorted on these carbon characteristics did not generally produce significant abnormal returns (alphas) after accounting for common risk factors like market movements, size, value, profitability, and investment, especially once exposure to carbon price (EUA) changes was included. Similarly, firm-level regressions (both Fama-MacBeth and panel OLS with fixed effects) did not show these carbon characteristics to be reliable predictors of returns on their own, after controlling for other known firm characteristics.

However, a key finding of this research is the evidence for conditional pricing of carbon risk. While the average relationship was weak, our Fama-MacBeth analysis revealed that the link between a firm's industry-adjusted emission levels (and intensity) and its stock returns significantly changes depending on movements in the European carbon market. Specifically, during months when EUA carbon prices increased, firms with higher industry-adjusted emissions or intensity tended to perform relatively better than their lower-emission peers. This statistically significant interaction suggests that the market's valuation of these carbon exposures is not static but rather depends on the prevailing carbon price environment, becoming more apparent when the cost of emitting carbon is actively rising. The direct economic size of this conditional return impact, however, appears modest on a monthly basis, suggesting its importance may lie more in understanding shifting risk perceptions than in identifying a large, exploitable return premium. This conditional effect was not found for the carbon intensity growth metric, nor were interactions with a broader Climate Policy Uncertainty index generally significant.

Finally, our investigation into whether a firm's direct sensitivity to carbon price changes (its "carbon beta") is itself a priced risk factor yielded inconclusive results. While simple portfolio sorts hinted at higher returns for firms with higher carbon betas, this relationship did not hold up in more rigorous multivariate Fama-MacBeth regressions that controlled for other firm characteristics.

In conclusion, this study of European equities from 2015 to 2024 indicates that simply holding firms with higher (or lower) industry-adjusted carbon emissions, intensity, or emissions growth did not consistently lead to statistically or economically significant abnormal returns on an unconditional basis. Instead, the pricing of carbon exposure related to emission levels and intensity appears to be more nuanced, becoming evident in a conditional manner that is linked to the dynamics of the carbon market itself. This highlights the importance of considering the prevailing economic and regulatory context when assessing the financial implications of carbon transition risk. While a clear, simple "carbon premium" or "greenium" was not consistently found for the tested metrics in an unconditional sense, the conditional effects suggest that carbon risk is indeed a factor that markets react to, particularly when its costs become more immediate.

Bibliography

Bolton, Patrick and Marcin Kacperczyk (Nov. 2021). “Do investors care about carbon risk?” In: *Journal of Financial Economics* 142.2, pp. 517–549. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2021.05.008.

— (Dec. 2023). “Global Pricing of Carbon-Transition Risk”. In: *Journal of Finance* 78.6, pp. 3677–3754. ISSN: 15406261. DOI: 10.1111/jofi.13272.

Choi, Darwin, Zhenyu Gao, and Wenxi Jiang (2020). “Source: The Review of Financial Studies”. In: 33.3, pp. 1112–1145. DOI: 10.2307/48574181. URL: https://www.jstor.org/stable/10.2307/48574181.

Dietz, Simon et al. (July 2016). “Climate value at risk’ of global financial assets”. In: *Nature Climate Change* 6.7, pp. 676–679. ISSN: 17586798. DOI: 10.1038/nclimate2972.

Energy Agency, International (2024). *World Energy Outlook 2024*. Tech. rep. URL: www.iea.org/terms.

European Comission (Apr. 2024). “Record reduction of 2023 ETS emissions due largely to boost in renewable energy - European Commission”. In: *European Comission*. URL: https://climate.ec.europa.eu/news-your-voice/news/record-reduction-2023-ets-emissions-due-largely-boost-renewable-energy-2024-04-03_en.

Fama, Eugene F and Kenneth R French (1993). *Common risk factors in the returns on stocks and bonds**. Tech. rep., pp. 3–56.

Gavriilidis, Konstantinos (n.d.). *Measuring Climate Policy Uncertainty*. Tech. rep.

Hong, Harrison and Marcin Kacperczyk (July 2009). “The price of sin: The effects of social norms on markets”. In: *Journal of Financial Economics* 93.1, pp. 15–36. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2008.09.001.

Hsu, Po Hsuan, Kai Li, and Chi Yang Tsou (June 2023). “The Pollution Premium”. In: *Journal of Finance* 78.3, pp. 1343–1392. ISSN: 15406261. DOI: 10.1111/jofi.13217.

Iii, Robert F Engle et al. (2019). *Hedging Climate Change News*. Tech. rep.

Ilhan, Emirhan, Zacharias Sautner, and Grigory Vilkov (2021). “Source: The Review of Financial Studies”. In: 34.3, pp. 1540–1571. DOI: 10.2307/48616797. URL: https://www.jstor.org/stable/10.2307/48616797?seq=1&cid=pdf-.

Krueger, Philipp, Zacharias Sautner, and Laura T Starks (2020). “The Importance of Climate Risks for Institutional Investors”. In: 33.3, pp. 1067–1111. DOI: 10.2307/48574180. URL: https://www.jstor.org/stable/10.2307/48574180.

McKinsey (2024). *What are Scope 1, 2, and 3 emissions?* Tech. rep.

Oestreich, A. Marcel and Ilias Tsiakas (Sept. 2015). "Carbon emissions and stock returns: Evidence from the EU Emissions Trading Scheme". In: *Journal of Banking and Finance* 58, pp. 294–308. ISSN: 03784266. DOI: 10.1016/j.jbankfin.2015.05.005.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor (Nov. 2021). "Sustainable investing in equilibrium". In: *Journal of Financial Economics* 142.2, pp. 550–571. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2020.12.011.

Riordan, Ryan (n.d.). *Carbon Risk**. Tech. rep. URL: https://ssrn.com/abstract=2930897.

Sharpe, William F. (1964). "CAPITAL ASSET PRICES: A THEORY OF MARKET EQUILIBRIUM UNDER CONDITIONS OF RISK". In: *The Journal of Finance* 19.3, pp. 425–442. ISSN: 15406261. DOI: 10.1111/j.1540-6261.1964.tb02865.x.

World Bank (2024). "State and Trends of Carbon Pricing". In: DOI: 10.1596/978-1-4648-2127-1. URL: www.worldbank.org.

Appendix A

Appendix

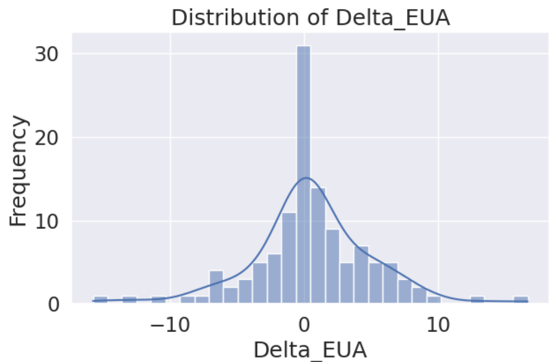

(a) Distribution of ΔEUA (monthly change in EU allowance price).

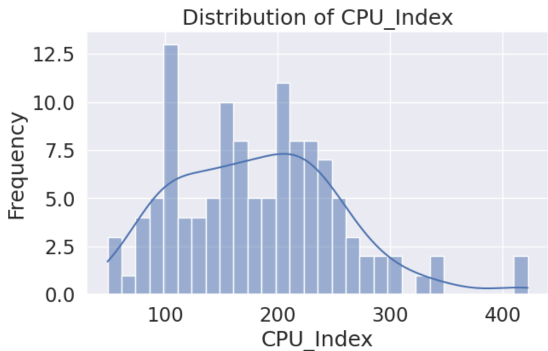

(b) Distribution of CPU Index (climate-policy uncertainty).

Figure A.1: Histograms of the two climate-risk factors.

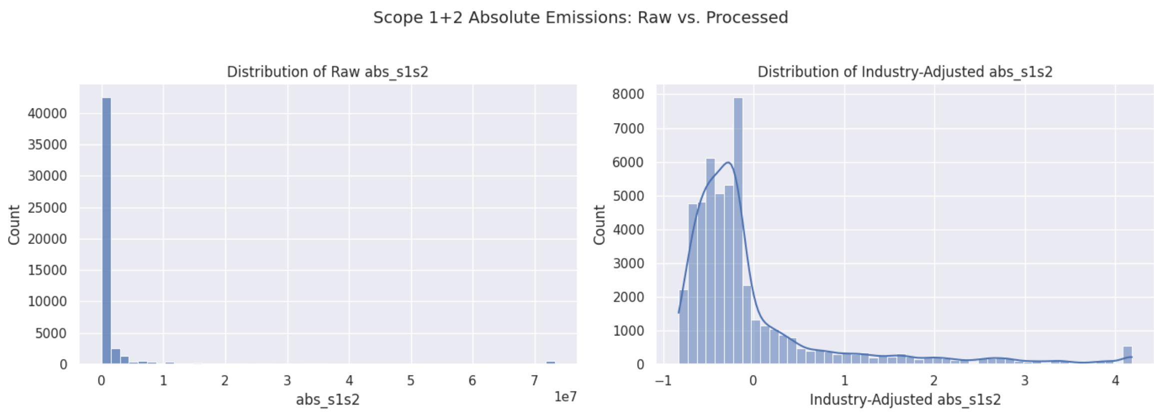

Figure A.2: Scope 1+2 absolute emissions: raw vs. industry-adjusted.

Table A.1: Variable Definitions and Construction

Panel A: Firm-level variables				
Variable (symbol)	**Definition / Formula**	**Units**	**Transformation**	**Data Source**
Excess return ($R_{i,t} - R_{f,t}$)	Monthly stock total return minus 1-month EURIBOR.	decimal	none	Refinitiv Eikon
Absolute emissions ($\text{ABS}_{i,t-1}$)	Scope 1 + 2 CO_2e, calendar year $t-1$.	tCO_2e	Winsorise p1/p99	Refinitiv Eikon
Carbon intensity ($CI_{i,t-1}$)	$\frac{\text{ABS}_{i,t-1}}{\text{Revenue}_{i,t-1}}$	tCO_2e / €mm	$\ln(\cdot)$, winsorise	Refinitiv Eikon
Industry-adj. intensity ($CI^{\text{adj}}_{i,t-1}$)	z-score of CI within ICB industry × year.	z-score	standardise	Author calc.
Growth in intensity ($\Delta CI^{\text{adj}}_{i,t-1}$)	$\ln(CI_{i,t-1}) - \ln(CI_{i,t-2})$, then industry z-score.	z-score	see def.	Author calc.
Size ($\ln \text{MktCap}_{i,t-1}$)	Market value at $t-1$.	€	ln	Refinitiv Eikon
Book-to-market ($\ln B/M_{i,t-1}$)	Book equity / market cap.		ln	Refinitiv Eikon
Momentum ($MOM_{i,t-1}$)	Cum. return $t-12$ to $t-2$.	decimal	none	Refinitiv Eikon
Profitability ($ROE_{i,t-1}$)	EBIT / book equity.	%	none	RefinitivEikon
Panel B: Market and climate factors				
$MKT-RF_t$	CRSP Europe value-weighted market return minus 1m EURIBOR	%	none	French Data Lib.
SMB_t	Size factor (small minus big)	%	none	French Data Lib.
HML_t	Value factor (high minus low B/M)	%	none	French Data Lib.
RMW_t	Profitability factor (robust minus weak)	%	none	French Data Lib.
CMA_t	Investment factor (conservative minus aggressive)	%	none	French Data Lib.
ΔEUA_t	Monthly change in EU ETS front contract	€/t	none	Factset
CPU_t	Climate-policy-uncertainty index	index	none	(Gavriilidis n.d.)

Table A.2: Sample Coverage by Year (2015–2024)

Year	**Firms**	**Firm–Month Obs.**
2015	536	6 364
2016	547	6 495
2017	552	6 599
2018	557	6 655
2019	564	6 718
2020	569	6 789
2021	579	6 901
2022	585	6 989
2023	589	7 036
2024	593	7 097
Average / Total	567.1	67 643

Table A.3: Industry Composition of the Sample (2015–2024)

ICB super-sector	**Share (%)**
Industrial Goods & Services	17.36
Health Care	9.14
Banks	8.42
Financial Services	5.59
Chemicals	5.26
Basic Resources	4.82
Utilities	4.43
Food & Beverage	4.30
Technology	4.05
Construction & Materials	3.88
Insurance	3.34
Automobiles & Parts	3.14
Retail	2.92
Energy	2.87
Real Estate	2.81
Telecommunications	2.49
Travel & Leisure	2.44
Personal & Household Goods	2.36
Media	2.23
Aerospace & Defence	1.93
Total	100.00

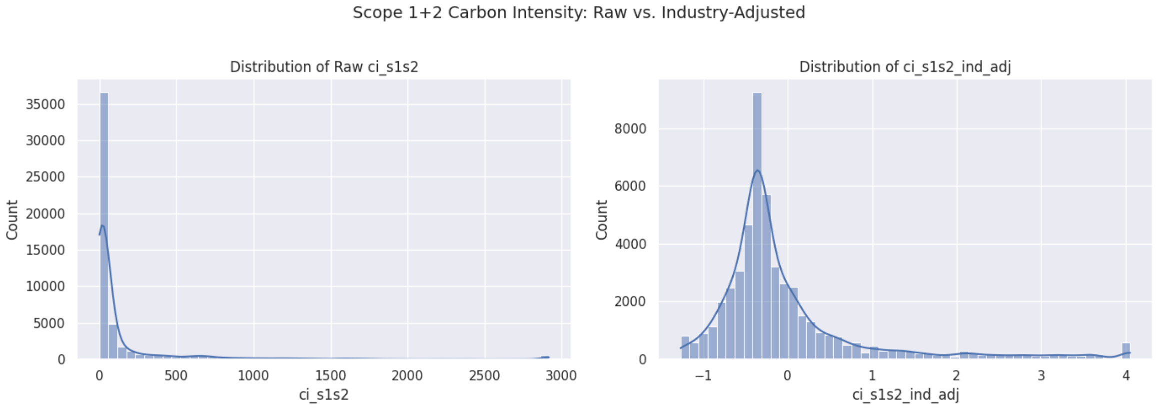

Figure A.3: Scope 1+2 carbon intensity: raw vs. industry-adjusted.

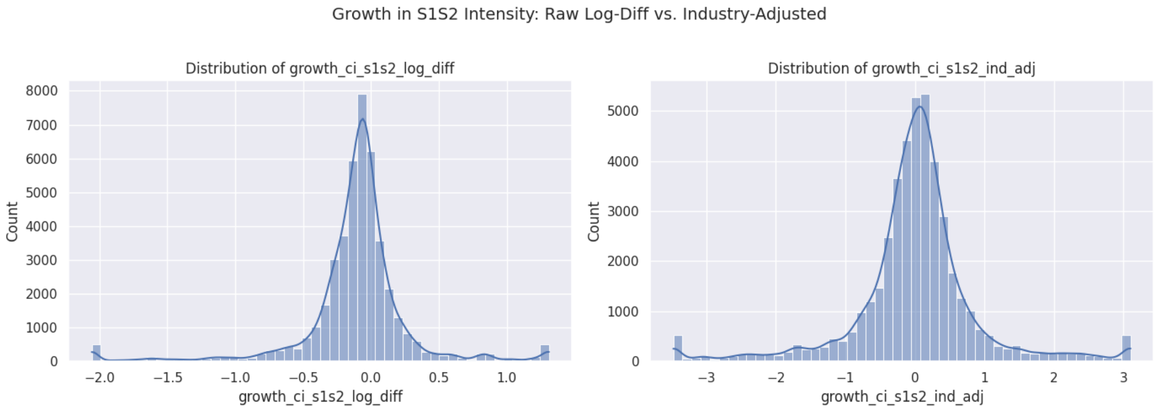

Figure A.4: Growth in intensity: raw log-diff vs. industry-adjusted.

Table A.4: Variance–Inflation Factors for Carbon Specifications

Variable	**Abs. Emissions**	**Carbon Intensity**	**Intensity Growth**
Carbon metric	1.36	1.02	1.00
log_bm_ratio	1.51	1.46	1.42
log_mkt_cap	1.34	1.02	1.03
momentum_12_1	1.03	1.03	1.02
profitability_metric	1.45	1.45	1.42
Maximum VIF	1.51	1.46	1.42

VIF is computed as $1/(1 - R_j^2)$, where R_j^2 is the R^2 from regressing variable j on all other regressors in the same specification. The highest VIF is 1.51, well below the conventional concern thresholds (5 or 10); multicollinearity is therefore not an issue.

Table A.5: ADF Stationarity Tests (AIC-selected lag length, monthly data)

Variable	N	**ADF stat (c)**	p**-value (c)**	**ADF stat (ct)**	p**-value (ct)**	**Conclusion @ 5%**
BMG Factor (AbsS1S2)	96	–4.517	0.0002	–8.583	0.0000	Stationary (c; ct)
BMG Factor (CiS1S2)	96	–11.101	0.0000	–11.045	0.0000	Stationary (c; ct)
BMG Factor (GrowthCiS1S2)	84	–8.911	0.0000	–8.948	0.0000	Stationary (c; ct)
CPU Index	120	–3.671	0.0045	–7.643	0.0000	Stationary (c; ct)
ΔEUA	120	–12.125	0.0000	–12.075	0.0000	Stationary (c; ct)
MKT–RF (excess)	120	–10.980	0.0000	–10.936	0.0000	Stationary (c; ct)
SMB	120	–10.860	0.0000	–11.406	0.0000	Stationary (c; ct)
HML	120	–10.080	0.0000	–10.236	0.0000	Stationary (c; ct)
RMW	120	–8.075	0.0000	–8.431	0.0000	Stationary (c; ct)
CMA	120	–3.646	0.0049	–3.703	0.0222	Stationary (c; ct)

Notes: ADF statistics are reported under two specifications: (c) constant only; (ct) constant plus deterministic trend. Lag length selected by AIC for each series. "Stationary (c; ct)" indicates the null of a unit root is rejected at the 5% level under both specifications.