The Price of Emissions: Carbon Risk in the European Equity Market*

Lewei He†[1,2], Harald Lohre[1,2], Ingmar Nolte[1], and Chelsea Yao[1]

[1]Lancaster University Management School, Lancaster, United Kingdom
[2]Robeco, Rotterdam, Netherlands

July 31, 2025

Abstract

We investigate whether carbon risk commands a return premium in the European equity market. Using data from the EU Emissions Trading System (EU ETS) between 2013 and 2024, we construct two complementary measures of carbon risk. The first is a forward-looking, market-based proxy derived from the one-year forward convenience yield in the carbon futures market, which reflects investor expectations about future allowance prices. The second is an accounting-based proxy calculated as the ratio of estimated carbon expenses to key income statement items for firms regulated under the EU ETS. Both measures reveal meaningful operational exposure to carbon expenses of companies. By 2023, firms in the top quartile of carbon expense intensity faced carbon expenses greater or equal to 12.4% of operating income. Our results highlight a divergence: the market-based proxy is associated with a significant risk premium in the broad STOXX Europe 600, but within the subsample of firms directly regulated by the EU ETS, neither the market-based nor the accounting-based proxy is associated with a return premium.

JEL classification: G12, Q41.

Keywords: Carbon Risk, EU ETS, Convenience Yield, Stock Returns.

---

*We thank Dong Lou and Laurens Swinkels for helpful comments and suggestions. We would like to acknowledge financial support from the Economic and Social Research Council (ESRC) North West Social Science Doctoral Training Partnership (NWSSDTP) CASE Studentship. Note that this paper expresses the authors' views which do not necessarily coincide with those of Robeco.

†Corresponding author, Email: l.he6@lancaster.ac.uk

1 Introduction

Climate policy imposes increasing financial constraints on firms, especially in carbon-intensive sectors. This raises a central question in asset pricing: do equity markets reward or penalize firms for their carbon exposure? In Europe, the primary mechanism for pricing carbon is the EU Emissions Trading System (EU ETS), a cap-and-trade program covering over 10,000 installations across the electricity, heat generation, industrial manufacturing, and aviation sectors, accounting for over 40% of European Union's emissions. By restricting the supply of carbon allowances and enabling trading, the EU ETS imposes a direct and measurable cost on emitting firms. While economic costs of compliance are clear, it remains an open question whether financial markets treat this exposure as a priced risk factor.

The financial exposure created by the EU ETS is one example of what is broadly referred to as carbon risk. This risk captures the potential costs companies may incur due to carbon emissions, whether through regulation, compliance, or reputational pressures. Carbon risk is often considered a component of transition risk, which reflects the broader economic impact of the shift toward a low-carbon economy. Firms with high carbon intensity may face higher operating costs, regulatory penalties, or asset impairments. While these risks have clear implications for valuation, existing evidence on whether they are priced in stock markets is mixed: some studies report that green stocks outperform brown stocks, others find the opposite, and still others detect no systematic difference. Eskildsen, Ibert, Jensen, and Pedersen (2024) attribute these divergent findings to differences in greenness definitions, data handling, and test design.

To address this ambiguity, we examine a narrower and more concrete dimension of climate exposure: the pricing of carbon risk in European equity returns. Two complementary measures are employed. The first is a forward-looking, market-based indicator derived from the one-year forward convenience yield in the carbon futures market, which reflects carbon market investor expectations about future allowance prices. The second is a firm-level accounting measure that captures the actual financial burden of carbon compliance, calculated as the ratio of carbon expenses to key income statement items. By examining both expectations and realized costs, we test whether either dimension of carbon risk is reflected in equity valuations.

The EU ETS was launched in 2005 and is currently in its fourth phase, which runs from 2021 to 2030. Our analysis focuses on Phases 3 and 4, covering the period from 2013 to 2024. The first two phases are excluded due to structural limitations. Phase 1 (2005 to 2007) was a pilot period with substantial overallocation of allowances, while Phase 2 (2008 to 2012) coincided with the global financial crisis which prompted major regulatory reforms that informed Phase 3. Hence, structural features of Phase 1 and 2 undermine the consistency and reliability of early carbon price signals, making Phase 3 a more appropriate starting point for studying asset pricing effects.

Our analysis focuses on firms in the STOXX Europe 600 index, including large-, mid-, and small-cap companies across 17 European countries. To measure market-based carbon risk, we examine pricing in the carbon futures market, where EU Allowances (EUAs) are actively traded on the Intercontinental Exchange (ICE). We begin by documenting a strong upward trend in EUA futures prices, which rose from under €10 in 2013 to nearly €70 by 2024. This rise reflects increasing regulatory stringency and a tightening carbon market, making futures prices a meaningful source of forward-looking carbon risk information.

The one-year forward convenience yield is calculated using the cost of carry model, substituting the spot and future with the prompt- (current) and front-year December contract due to time difference considerations. We find that the annualised convenience yield remains persistently below –1% throughout the sample period. This negative yield is a counterintuitive result, but it is consistent with findings from earlier studies on the EU ETS by Trück and Weron (2016), Bredin and Parsons (2016), and Palao and Pardo (2021). In standard commodity markets, the convenience yield represents the benefit of holding the underlying asset and is expected to be positive by construction. In the context of the EU ETS, we interpret the persistently negative yield as an indicator that market participants expect carbon allowance prices to rise over the following year. This forward-looking pattern motivates our use of the return in convenience yield can serve as a market-implied signal of carbon risk, reflecting investor expectations of a more stringent regulatory environment and higher future compliance costs for emitting firms.

To test whether this market-implied carbon risk is reflected in equity prices, we sort stocks on their sensitivity to our proxy on a rolling basis with 3-year estimation period and 1-month holding period. We find that this risk is priced in the broad STOXX Europe 600 sample, generating a significant return premium that persists after controlling for standard risk factors, with mean monthly return of 0.51% and a monthly alpha of 0.77%. In contrast, when we focus on the subsample of firms with direct EU ETS exposure, we find this risk premium disappears.

We also examine an accounting-based measure of carbon risk by estimating firm-level carbon expense ratios under the EU ETS. Using installation-level data from the Union Registry, we aggregate verified emissions to the firm level and multiply them by the average annual spot price of allowances. These carbon expenses are then scaled by net revenue, gross income, and operating income to assess their financial materiality. In 2023, firms in the top quartile of carbon expense intensity faced carbon expenses greater than or equal to 1.4% of revenue, 5.1% of gross income, and 12.4% of operating income. In heavily regulated sectors such as Materials, these ratios were substantially higher, reaching 5%, 17%, and 50%, respectively. These magnitudes suggest that for many firms, carbon compliance imposes a non-trivial drag on profitability.

To test whether these realized carbon expense ratios are reflected in stock prices, we conduct portfolio sorts on the firms with direct EU ETS exposure. Our results consistently show no evidence of a return premium. The return spread between firms with the highest and lowest carbon expense intensity is statistically insignificant, a finding that is robust across different expense measures, after adjusting for sector-neutrality, and in multi-factor models that yield insignificant alphas. This lack of pricing also persists in double-sorting tests that account for firms' sensitivity to our market-based risk proxy. Together, these findings indicate that carbon expense intensity, despite its financial relevance, is not priced in the European equity market.

One key contribution of this study is the proposal of two distinct measures of carbon risk: a market-based indicator, the carbon convenience yield, and an accounting-based metric, the carbon expense ratio. As the field of climate finance expands, a variety of climate risk proxies have been developed. As summarised in Table 1 of Eskildsen et al. (2024), the most commonly used "greenness" measures in academic research are sourced from Trucost, MSCI, Sustainalytics, and the Toxics Release Inventory (TRI). Trucost, used by studies such as Bolton and Kacperczyk (2021), Aswani, Raghunandan, and Rajgopal (2024), Zhang (2025), Ardia, Bluteau, Boudt, and Inghelbrecht (2023), and Matsumura, Prakash, and Vera-Munoz (2014), is published by S&P and largely based on company-reported Scope 1 and Scope 2 emissions, supplemented by proprietary estimation models. MSCI and Sustainalytics ESG scores, used in Engle, Giglio, Kelly, Lee, and Stroebel (2020) and Pedersen, Fitzgibbons, and Pomorski (2021), are ratings produced by commercial providers. TRI, used in Hsu, Li, and Tsou (2023) and Akey and Appel (2021), is maintained by the U.S. Environmental Protection Agency and focuses on toxic chemical releases and pollution-prevention activities. Except for TRI, most greenness data originate from for-profit third parties with complex methodologies, which may introduce biases (Berg, Kölbel, and Rigobon, 2022) and raise concerns about potential washing (Berg, Fabisik, and Sautner, 2021).

Another stream of literature develops text-based measures to capture climate risk. For instance, Engle et al. (2020) construct the Wall Street Journal Climate Change News Index to capture media attention, while Sautner, Van Lent, Vilkov, and Zhang (2023) extract climate-related keywords from corporate earnings calls to quantify firm-specific exposure. Although such approaches offer valuable insights, they face limitations. News-based indicators may be influenced by the "news value" principle (Galtung and Ruge, 1965; Harcup and and O'Neill, 2001), which biases coverage toward acute, event-driven issues and underrepresents slowmoving but persistent risks such as climate change. At the same time, reliance on company presentations and more broadly self-reporting remains vulnerable to reporting inconsistencies and greenwashing. In contrast, our two proposed measures are grounded in observable market activity and audited operational data. The carbon convenience yield is derived from carbon futures prices, capturing forward-looking investor expectations, while the carbon expense ratio reflects realised compliance costs using verified EU ETS data. Together, these measures offer a more direct and transparent lens on how carbon risk is both anticipated and borne by firms in financial markets. Moreover, they explicitly address both policy & legal and market transition risks by identifying carbon expenditures within the income statement, thereby shedding light on the financial impact of carbon exposure, as outlined by Task Force on Climate-Related Financial Disclosures (2017).

Figure 1. Climate-Related Risks, Opportunities, and Financial Impact. Source: TCFD Recommendations.

A closely related study by Fuchs, Stroebel, and Terstegge (2024) develop a Carbon VIX from options prices on carbon allowances. They show that the stock returns of carbon solution providers vary positively with carbon prices and negatively with carbon price uncertainty, as proxied by the Carbon VIX, providing another example of a market-implied measure of carbon risk.

The key distinction between the carbon convenience yield and the Carbon VIX lies in both the type of information each embeds and the structure of the underlying market. First, the carbon convenience yield captures directional information, the widening or narrowing of the spread between prompt and front futures, whereas the Carbon VIX measures variance without directional signal. Second, the market participation base differs. While options volumes have grown, representing 27% of total notional in December 2021 according to Figure 21 of ESMA (2022), futures still dominate at 53%. Moreover, the holder composition varies substantially. As shown in Figure 56 of ESMA (2022), nearly 100% of volumes held by compliance entities and other non-financial participants are futures, whereas options are predominantly held by investment firms, credit institutions, and funds. This difference in market structure suggests that the futures market reflects a broader spectrum of market expectations, especially from the actual "consumers" of allowances, that is, the firms subject to compliance obligations. Consequently, the carbon convenience yield may offer a more representative and policy-relevant signal of transition risk in addition to volatility-based measures derived from financial intermediaries' trading.

This study also contributes to the growing debate on whether carbon risk is priced in the stock market. Hsu et al. (2023) argue that, in theory, brown stocks should offer a pollution premium due to their greater exposure to transition risk. However, empirical results have been mixed. A key study by Bolton and Kacperczyk (2021) finds that firms' absolute emissions are reflected in stock returns, whereas carbon intensity is not. This finding has been challenged by more recent work. Aswani et al. (2024) show that the results in Bolton and Kacperczyk (2021) are driven by vendor-estimated emissions rather than firm-disclosed data, and argue that carbon intensity is a more appropriate measure of carbon exposure. Zhang (2025) further finds that the documented carbon premium may stem from data release lags, and once those are accounted for, carbon risk appears to carry a negative premium in U.S. markets. Building on this debate, our study introduces two alternative measures of carbon risk: the carbon convenience yield return and carbon expense ratios, the latter of which can also be interpreted as intensity-based. We find evidence that the first proxy is priced in the broad STOXX Europe 600 universe, while neither is priced in the group with direct exposure to EU ETS.

This study also contributes to the literature on the EU ETS and the pricing of carbon allowances. One major strand of this literature examines the determinants and dynamics of carbon prices. For example, Mansanet-Bataller, Pardo, and Valor (2007) use data from the first year of Phase 1 (2005) and show that changes in Brent crude oil and natural gas prices, as well as extreme temperatures, influence carbon prices. Alberola, Chevallier, and Chèze (2008) extend the analysis to the full Phase 1 (2005–2007) period, confirming these findings and identifying two structural breaks that altered the fundamental drivers of carbon prices. Chevallier (2009) study macroeconomic risk factors and report that neither T-bill yields nor excess returns on a global commodity portfolio significantly affect carbon futures returns, suggesting that participants in the carbon market behave differently from those in broader financial markets. More recently, Batten, Maddox, and Young (2021) focus on Phase 3 and find that oil, coal, and electricity prices, along with the clean spark spread, jointly explain variation in carbon prices. Ye and Xue (2021) propose a dictionary-based carbon tone index and show that it has predictive power for carbon price returns. Another line of research seeks to improve the modelling of carbon price dynamics. For example, Lutz, Pigorsch, and Rotfuß (2013) apply a Markov regime-switching GARCH model to capture shifts in the data-generating process during economic shocks. Zhu, Ye, Han, Wang, He, Wei, and Xie (2019) develop a multiscale decomposition model that groups carbon price drivers into economically meaningful categories, improving in-sample fit.

A separate branch of the literature explores the role of EUAs as an asset class and their use in portfolio management. Medina and Pardo (2013) examine stylised facts of EUA returns and find that they display characteristics of both commodity and financial futures, suggesting that EUAs may represent a distinct asset class. Afonin, Bredin, Cuthbertson, Muckley, and Nitzsche (2018) provide evidence of portfolio diversification benefits from short positions in EUA futures during Phase 1, but not in Phase 2. Demiralay, Gencer, and Bayraci (2022) document time-varying hedging and diversification benefits of EUA futures, although these are less effective than those of precious metals or agricultural commodities and deteriorated during the COVID-19 pandemic. Expanding the scope beyond the EU ETS, Azlen, Gostlow, and Child (2022) and Swinkels and Yang (2023) study global carbon markets and propose that limited return correlation across jurisdictions offers diversification potential. They further argue that a global carbon composite index can serve as a hedge for conventional asset classes.

Most relevant to this paper, the persistence of a negative carbon convenience yield has been documented by Trück and Weron (2016), Bredin and Parsons (2016), and Palao and Pardo (2021), and is discussed in greater detail in Section 4.2. We confirm this empirical pattern using data from Phase 4 of the EU ETS. However, rather than modelling the drivers of the convenience yield itself, we reinterpret it as a forward-looking proxy for carbon transition risk and test whether this market-based signal is priced in the cross-section of European stock returns.

The remainder of the paper is structured as follows. Section 2 provides background on the EU Emissions Trading System. Section 3 describes the data sources, including EUA futures data, installation- and firm-level emissions records, and financial statement items. Section 4 introduces the carbon convenience yield, detailing its construction and examining its role as a market-implied signal of carbon risk. Section 5 presents the carbon expense ratio as a firm-level accounting-based measure and analyses its financial relevance. Section 6 concludes the paper.

2 EU ETS

The EU Emissions Trading System (EU ETS) was established as a policy response to the Kyoto Protocol adopted in 1997 to meet emissions reduction targets. The concept was first introduced by the European Commission in 2000, with the Directive formally adopted in 2003. Launched in 2005, the EU ETS is the first and currently the largest carbon market in the world. It currently covers emissions from the electricity and heat generation, industrial manufacturing, and aviation sectors, with around 10,000 installations, accounting for approximately 40% of total greenhouse gas (GHG) emissions in the EU. The EU ETS operates under a cap and trade principle, where a limit is set on the total amount of GHG that can be emitted by all installations covered by the system. This cap is reduced annually in line with the EU's climate targets.

Participating installations are required to monitor and report their annual emissions, which must be matched by surrendering an equivalent amount of allowances. If the surrendered allowances fall short of actual emissions, substantial fines apply, currently set at €100 per excess tonne. Allowances are primarily distributed through auctions, although some companies receive free allocations, for sectors with the highest risk of relocating their production outside of the EU, and to modernise the energy sector. These allowances can be traded on the market, allowing firms that reduce emissions to sell or bank surplus allowances.

This mechanism ensures that polluters bear the cost of their greenhouse gas emissions, while providing financial incentives for emission reductions. Revenues generated from allowance auctions are allocated to EU Member States to support green transition initiatives, creating a reinforcing cycle aimed at achieving climate neutrality by 2050.

After its establishment, the EU ETS has evolved through several phases, each marked by changes in coverage scope, allocation methods, and regulatory specifications. Phase 1 (2005–2007) served as a pilot phase focused on developing the necessary infrastructure for an operational trading system. During this phase, only power generators and energy-intensive industries were included. However, because the emissions cap was set based on estimations without verified data, the supply of allowances far exceeded demand. Combined with the near-universal free allocation of allowances, this oversupply led to a collapse in carbon prices, which dropped to zero in 2007.

Phase 2 (2008–2012) aligned with the first commitment period of the Kyoto Protocol. With verified emissions data from Phase 1, the cap was tightened, and a portion of allowances began to be auctioned instead of entirely allocated for free. Despite this progress, the global financial crisis in 2008 caused a significant drop in industrial activity, resulting in a persistent surplus of allowances. Although the fine for non-compliance was raised to €100 per excess tonne of emissions, the oversupply continued to depress carbon prices.

Phase 3 (2013–2020) introduced substantial reforms. Auctions became the primary method for distributing allowances, replacing the previous default of free allocation. The scope of the system expanded to include additional sectors and gases, improving market depth and coverage. The Market Stability Reserve (MSR) was introduced to address the structural surplus of allowances that had accumulated in earlier phases, helping to support price stability.

Phase 4 (2021–2030) largely continues the regulatory structure introduced in Phase 3, with a progressively declining cap to meet the EU's climate targets. In this phase, up to 57% of allowances are distributed via auction, while the remainder are allocated for free, consistent with the goals outlined in the Fit for 55 legislative package, as to reduce EU GHG emission by at least 55% comparing with 1990 levels by 2030.[1]

As the first and largest carbon market in the world, covering the entire EU area, the EU ETS provides an important framework for understanding carbon pricing and its implications for firm operations and financial markets. Since both the regulatory structure and the price dynamics underwent significant reforms between Phase 2 and Phase 3, and given that the current framework is expected to remain in place for the foreseeable future, this study focuses on the period from the beginning of Phase 3 in 2013 through to the end of 2024.

[1] For Fit for 55 see: `https://www.consilium.europa.eu/en/policies/fit-for-55/`

3 Data

3.1 EUA Futures

The data related to the carbon market are sourced from multiple platforms. The secondary market for EU Allowances (EUAs) is primarily facilitated through two exchanges: ICE Endex (ICE) and the European Energy Exchange (EEX). Both platforms offer a range of futures contracts, including daily futures, monthly futures for the current and next two months, quarterly futures, and December futures. Among these, the ICE daily futures (treated as spot prices due to their daily expiry) and the ICE December futures are used in this study, primarily due to their superior liquidity.

Figure 2 presents the trading volumes of these contracts in continuation form, where each series automatically rolls over to the next available contract upon expiry. The data clearly show that ICE December futures consistently exhibit the highest trading volume, compared with both ICE and EEX monthly contracts. Notably, the spikes in volume for monthly continuations toward year-end correspond to the rollover from November to December contracts. The figure also highlights ICE's dominant position in the EUA market compared to EEX. According to ESMA (2022), "*ICE Endex has a dominant position with 85% of the total positions while EEX holds the remaining 15%*".

Despite the disparity in trading volumes, the price differential between exchanges has remained close to zero since mid-2015, with only nine instances exceeding €0.05 over the period. Therefore, using ICE December futures in this analysis does not introduce any material bias. EUA futures data are retrieved from Refinitiv Eikon.

3.2 Euro Area Interest Rate

In constructing the convenience yield, an appropriate forward-looking risk-free rate is required. Unlike the United States and other countries where Treasury securities, issued by national governments and backed by full faith and credit, are readily available and reliable

Figure 2. ETS Secondary Market Trade Volume. This figure compares the trade volumes of three types of EUA futures contracts over the period from 2013 to 2024: the December contract traded on ICE (ICE December Continuation), the monthly contract traded on ICE (ICE Monthly Continuation), and the monthly contract traded on EEX (EEX Monthly Continuation).

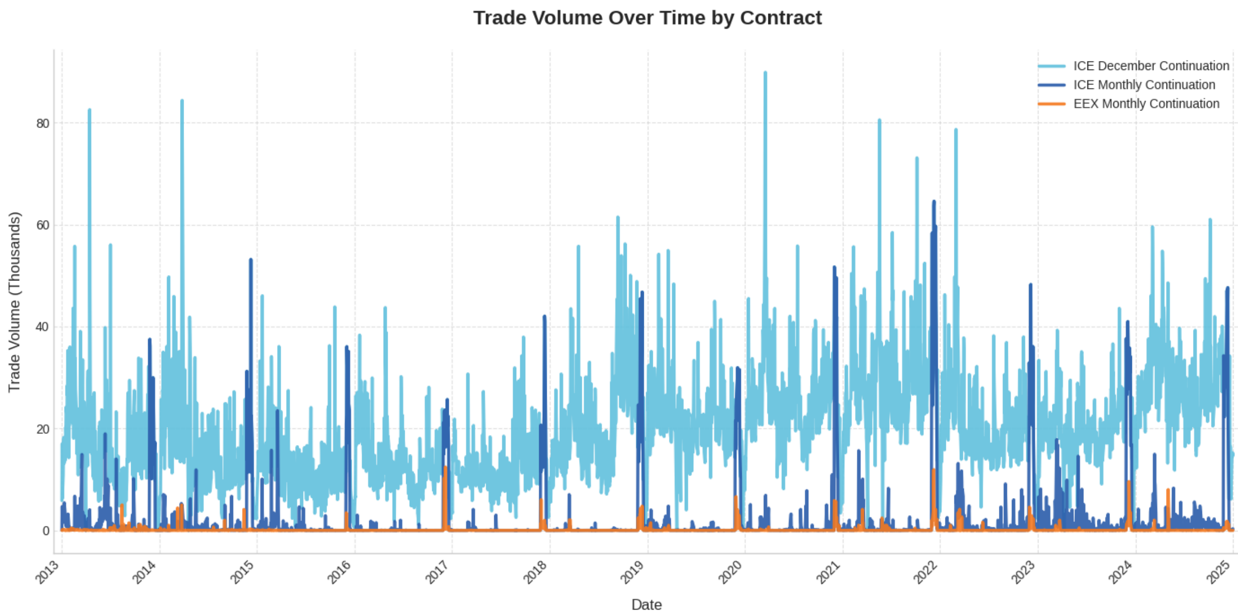

proxies, the European Commission only began issuing EU-Bills in September 2021, making them unsuitable for the full sample period of this study.[2]

Furthermore, because the futures contracts used in this analysis have time to expiration spanning months or years, short-term overnight rates such as the Euro Overnight Index Average (EONIA) and its successor, the euro short-term rate (€STR), are not adequate substitutes. Although German Bunds are often used as a euro area benchmark, it seems inappropriate to solely rely on German Bunds given the wider geographic coverage of the EU ETS, which includes multiple member states.

To address this, our study employs Euro area yield curves for AAA-rated bonds, as estimated by the European Central Bank (ECB). These curves are based on central government bonds with AAA ratings across the euro area and are published daily on the ECB's official

[2]For EU funding instruments see: `https://commission.europa.eu/strategy-and-policy/eu-budget/eu-borrower-investor-relations/funding-instruments_en`

website.[3]

3.3 EU ETS Installations

All EU ETS installation-level data are recorded in the Union Registry (UR), which has replaced the EU Transaction Log (EUTL) as of 2025. An installation refers to a specific production site or facility, a factory for example, that directly emits greenhouse gases. Each installation is linked to an operator account, which is held by the company directly responsible for the site. Verified emissions are reported, and allowance transactions (including holdings and surrenders) are conducted at the operator account level.

On the UR, the relationship between installations and operator account holders is explicitly defined, and a single account holder can be linked to multiple installations. A key challenge for this study is mapping these account holders to publicly listed parent companies, as such mappings are not readily disclosed in the registry. Since asset pricing research requires emissions exposure to be aggregated at the listed company level, a two-step matching procedure is used. First, a fuzzy match algorithm is applied to link operator account names in UR to company names in Refinitiv. Second, all candidate matches with a similarity score above 80 (100 stands for exact match) are manually reviewed to ensure accuracy and eliminate false positives. This process yields a verified mapping between operator accounts and publicly listed firms, enabling firm-level estimation of carbon exposure.

3.4 Stock Universe

The STOXX Europe 600 is the stock universe of choice for this study. Comprising 600 constituents, this index includes small-, mid-, and large-cap securities across 17 European countries and 11 industries, representing approximately 90% of the free-float market capitalization in the region.[4] It is widely regarded as the key benchmark for the European

[3] For Euro area yield curves see: `https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/euro_area_yield_curves/html/index.en.html`

[4] For STOXX Europe 600 see: `https://stoxx.com/index/sxxp/`

stock market. A broader index such as the STOXX Europe Total Market, or an even more expansive universe, is not adopted due to the prevalence of multiple listings across European exchanges. Such duplication would introduce noise and ambiguity in matching publicly listed firms to EU ETS account holders, as discussed in Section 3.3.

The EU ETS does not apply uniformly across all firms or sectors in Europe. Instead, it targets emissions from specific activities, primarily electricity and heat generation, industrial manufacturing, and aviation. As a result, the sector composition of STOXX Europe 600 constituents with EU ETS operator accounts differs structurally from that of the full index.

Figure 3 compares the sector distribution between the full STOXX Europe 600 and the subset of firms with ETS operator accounts. The ETS-exposed subset is notably concentrated in manufacturing-related sectors such as Materials, Utilities, and Energy, which together represent over 41% of the sample, compared to less than 17% in the full index. This concentration is balanced by a reduction in the representation of service-oriented sectors like Financials, Information Technology, Real Estate, and Communication Services, which collectively decrease from 35% in the full sample to only 10% in the ETS-exposed group.

This shift in sector composition has two implications for the analysis that follows. First, it provides context for the sector-level carbon expense ratios presented in Section 5.1. Second, it motivates the use of sector-neutral portfolio sorts in Section 5.2 to ensure that return patterns are not driven by industry effects.

Monthly return, market capitalisation, yearly net revenue, gross income, and operating income data are obtained from Refinitiv Datastream. Value-weighted European Fama-French factors are retrieved at a monthly frequency from the Kenneth French Data Library.[5]

[5]Kenneth French Data Library: `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`

Figure 3. Sector Weights and Deviation for full STOXX Europe 600 sample and the subset with with involvement in the EU ETS. This figure compares the sectoral distribution of firms in the full STOXX Europe 600 sample and the subset with involvement in the EU ETS. The sectors are sorted by the weight difference between the subsample with ETS exposure and the full sample.

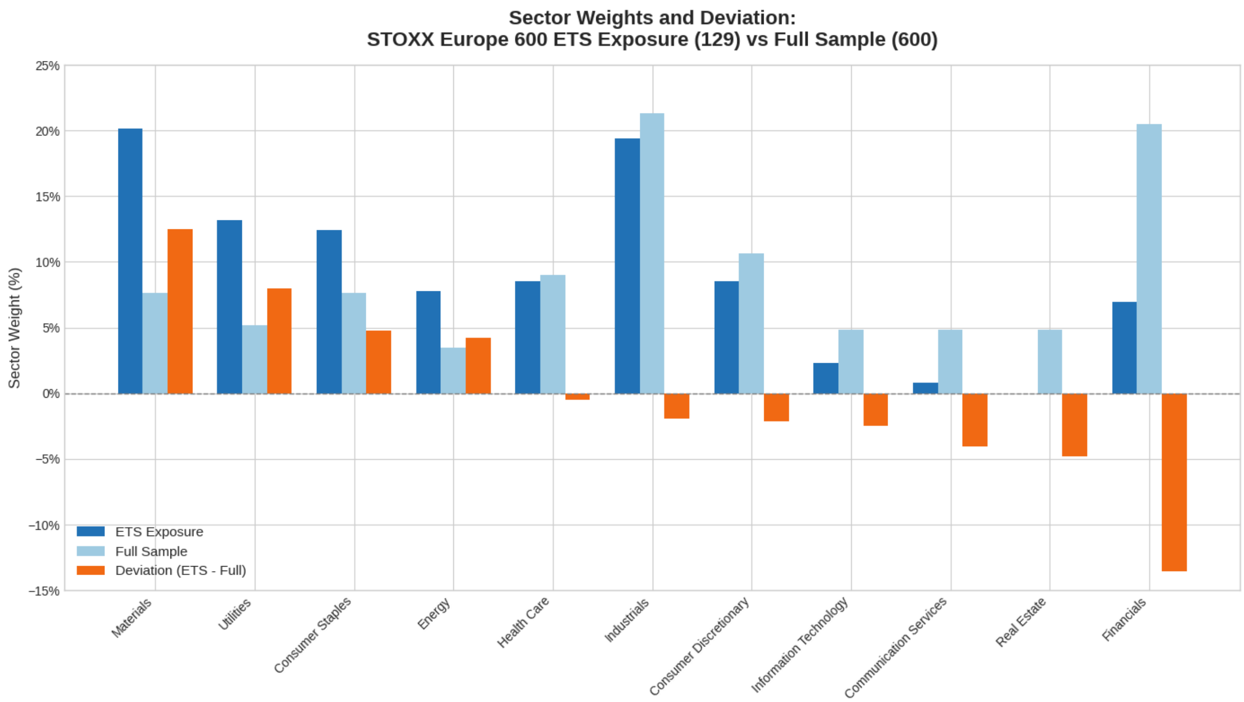

4 Carbon Price and Carbon Convenience Yield

4.1 Carbon Price

Figure 4 plots the ICE EUA Daily Futures price over the period from 2013 to 2024. The trend can be broadly divided into several distinct phases. From 2013 through the end of 2017, prices remained at single-digit levels, possibly reflecting a persistent market surplus following the overallocation of allowances in earlier phases of the EU ETS. A notable upward shift began in 2018, coinciding with the establishment of the Market Stability Reserve (MSR), which was designed to absorb excess allowances and may have contributed to rising price expectations. Prices doubled to exceed €20 per tonne and remained relatively stable between €20 and €30 until the onset of the COVID-19 pandemic in early 2020, when a temporary drop in industrial activity likely contributed to a brief decline below €20. A sharp rally followed, with prices reaching approximately €95 by early 2022, an increase of nearly fivefold in two years. This surge may have been influenced by strengthened EU climate policy initiatives, such as the “Fit for 55” package, as well as the 2021 to 2022 energy crisis, which saw elevated gas prices and a shift toward coal consumption, potentially increasing short-term EUA demand. Notably, under current regulation, EUA prices are unlikely to exceed €100 per tonne, as firms face a fixed penalty of €100 per tonne of excess emissions without allowance coverage. After a year of heightened volatility, the market entered a descending channel in early 2023, declining from €95 to around €50. Prices partially recovered in 2024, generally trading in a range between €60 and €70.

Figure 4. ICE EUA Daily Futures Price. This figure displays the trend in the ICE EUA Daily Futures price, commonly considered a proxy for the spot price of EU carbon allowances, over the period from 2013 to 2024.

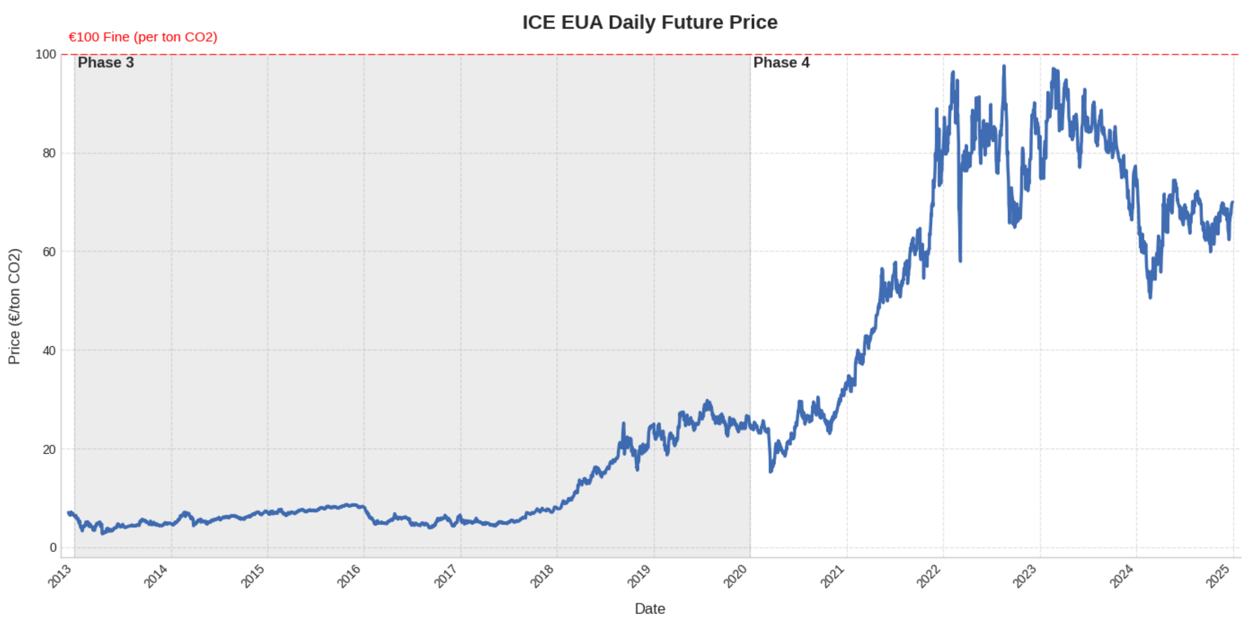

4.2 Carbon Convenience Yield

The cost of carry model is a standard approach for linking spot and futures prices in commodity markets. It relates the futures price to the spot price by incorporating the opportunity cost of capital and the benefits and costs of holding the underlying asset. The relationship is given by:

$$F_{t,T} = S_t \times e^{(r_{t,T} + sc_{t,T} - cy_{t,T})(T-t)},$$

where $F_{t,T}$ is the futures price on day t, with maturity date T; S_t is the spot price on day t; $r_{t,T}$ is the risk-free rate over the interval from t to T; $sc_{t,T}$ is the storage cost over that period; and $cy_{t,T}$ is the convenience yield.

In traditional commodity markets, the convenience yield represents the non-monetary benefits of physically holding the asset, such as inventory flexibility and protection against supply disruptions. These benefits compensate for storage costs, and the convenience yield is typically positive by definition. However, carbon allowances are not physical goods. They do not incur storage or transportation costs, and the cost is generally treated as zero, simplifying this model to:

$$F_{t,T} = S_t \times e^{(r_{t,T} - cy_{t,T})(T-t)}.$$

Three structural features of the EU ETS are worth considering when examining the relationship between spot and futures carbon prices, and in particular, the expected sign of the convenience yield.

First, the system operates on a fixed annual compliance cycle: regulated firms are required to surrender allowances in April each year to cover verified emissions from the previous calendar year. As a result, compliance demand is concentrated in the spring, but allowances acquired at any point before April satisfy the obligation equally.

Second, the annual cap, i.e. the total quantity of allowances issued, is predetermined before the start of each ETS phase and is known years in advance. Adjustments to the cap through the MSR mechanism are also rule-based and transparent, determined with the reference point of Total Number of Allowances in Circulation (TNAC) typically published each May. Moreover, auction schedules are normally available before August, both to take effect the following year. This transparency and advance planning reduce policy uncertainty around short-term supply.

Third, banking of allowances is unrestricted after 2013: EUAs held in operator accounts do not expire and may be carried forward indefinitely. Although banking was in place from 2008, the transition from Phase 2 to Phase 3 did not allow carryover; this constraint no longer applies since 2013.

Taken together, these institutional features imply that there is limited additional value to immediate access to allowances. Production plans are determined by emitters themselves, compliance deadlines are predictable, short-term supply shocks are unlikely to happen, and the value of allowances does not deteriorate over time. Hence, one may hypothesize that in an efficient market, futures prices should exceed spot prices only by the time value of money. Under this condition, the convenience yield should be close to zero.

Nevertheless, there are plausible conditions under which the convenience yield could be positive. For instance, a temporary mismatch between near-term compliance demand and short-term supply availability could arise, particularly in the spring, as firms finalize verified emissions and prepare to surrender allowances towards March and April. This seasonal concentration of compliance activity may create localized scarcity, increasing the value of immediate access to allowances.

Another theoretical justification for a positive convenience yield lies in regulatory risk. If market participants perceive a non-negligible probability of unexpected supply restriction, such as unanticipated MSR tightening or politically driven intervention, they may assign a premium to holding allowances today rather than relying on future delivery. In this case, a positive convenience yield could reflect a risk premium for short-term policy surprises.

However, such conditions have not been observed in the history of the EU ETS.

A third hypothesis considers the possibility of a negative convenience yield, which is counterintuitive in the traditional commodity framework, where the convenience yield is typically positive. However, the carbon market differs fundamentally from physical commodity markets. Allowances are fully bankable, without need to be stored in the traditional sense, and not tied to production. As such, the concept of "convenience" does not relate to inventory or operational needs, but rather to the market's expectations of future regulatory conditions.

In this context, a negative convenience yield reflects the expectation that allowance prices will rise over time, due to ongoing cap reductions, declining free allocations, and tightening climate policy. It signals anticipated future scarcity, not current shortage. The convenience yield therefore reflects the slope of the carbon futures curve, and a negative value indicates that forward prices exceed spot prices by more than what can be explained by the time value of money alone.

Since the future price $F_{t,T}$, spot price S_t, and risk-free rate $r_{t,T}$ can be directly observed or implied from the market, the convenience yield can be expressed as:

$$cy_{t,T} = r_{t,T} - \frac{\ln(F_{t,T}) - \ln(S_t)}{T - t}.$$

In practice, a key challenge is the changing time to maturity between spot and futures contracts. Although monthly contracts exist, trading is concentrated in December deliveries (Figure 2). As the December contract approaches expiry, the time horizon shrinks, then abruptly jumps to nearly one year after rollover, creating inconsistency. To address this, we compute the forward convenience yield using two consecutive December contracts, maintaining a constant one-year maturity difference:

$$cy_{t,(T_1,T_2)} = r_{t,(T_1,T_2)} - \frac{\ln(F_{t,T_2}) - \ln(F_{t,T_1})}{T_2 - T_1},$$

where T_1 is the expiration date of the prompt-year December contract (the contract that expires in the current year), and T_2 is the expiration date of the forward-year December contract (the contract that expires in future years).

The forward interest rate $r_t(T_1, T_2)$ is derived from the ECB's euro area yield curve using standard bond calculations:

$$(1 + r_{t,T_1})^{T_1 - t}(1 + r_t(T_1, T_2))^{(T_2 - T_1)} = (1 + r_{t,T_2})^{T_2 - t},$$

which rearranges to:

$$r_{t,(T_1,T_2)} = \left(\frac{(1 + r_{t,T_2})^{T_2 - t}}{(1 + r_{t,T_1})^{T_1 - t}} \right)^{1/(T_2 - T_1)} - 1.$$

Figure 5 shows that the one-year forward convenience yield is consistently negative over the sample period, often below –1% on an annualised basis. This is economically meaningful, especially during the long period when euro area interest rates were near zero or negative. Although the result may appear counterintuitive from a standard cost-of-carry perspective, it is not unprecedented.

Bredin and Parsons (2016) document a shift from positive to negative convenience yields in the EU ETS beginning in 2008 and focus on the evolving term structure of the yield. Trück and Weron (2016) find that the convenience yield is positively related to interest rates, and negatively related to inventory levels and carbon market volatility, supporting the interpretation of the convenience yield as a risk premium compensating for EUA price uncertainty. More recently, Palao and Pardo (2021) use an extended sample to show that financial market variables such as equity indices, bond yields, and volatility indices help explain movements in the convenience yield, suggesting a growing role for portfolio investors taking long positions in EUA futures and pushing up forward prices.

These findings reinforce the empirical validity of the negative convenience yield and support our interpretation that it reflects a market-implied expectation of rising future compliance costs. We therefore view the negative convenience yield as consistent with forward-

Figure 5. One Year Forward Convenience Yield. This figure shows the trend in the one-year forward convenience yield for EU carbon allowances over the period from 2013 to 2024.

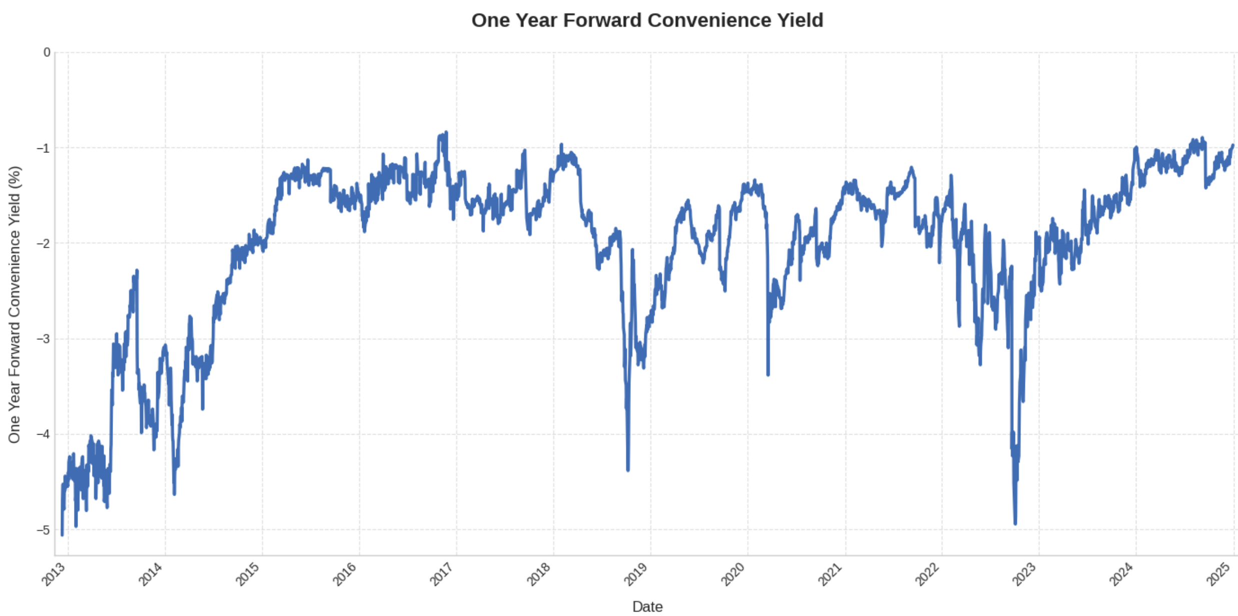

looking pricing of carbon transition risk.

While the level of the convenience yield reflects expectations, asset pricing theory suggests that expected returns are determined by exposures to innovations in systematic risks, not levels themselves. This is analogous to the Capital Asset Pricing Model (CAPM), where investors are not compensated for the level of the market, but for exposure to unexpected changes in aggregate returns. Similarly, if the convenience yield embeds expectations of future carbon regulation, it is the unexpected revisions to those expectations, captured by changes in the yield, that may be priced in stock returns.

Compared to changes in spot or futures prices, which may reflect short-term demand, macro volatility, or risk premia, changes in the convenience yield isolate relative shifts in regulatory expectations. This makes the change in the forward convenience yield a cleaner and theoretically consistent proxy for carbon transition risk shocks.

We now formalize this interpretation by examining the structure of changes in the forward convenience yield. Define:

$$\Delta cy_{t,(T_1,T_2)} = cy_{t,(T_1,T_2)} - cy_{t-1,(T_1,T_2)}.$$

Using a first-order Taylor expansion:

$$\Delta cy_{t,(T_1,T_2)} \approx \frac{\partial cy}{\partial F_{T_2}} \Delta F_{T_2} + \frac{\partial cy}{\partial F_{T_1}} \Delta F_{T_1} + \frac{\partial cy}{\partial r} \Delta r,$$

where the partial derivatives:

$$\frac{\partial cy}{\partial F_{T_2}} = -\frac{1}{(T_2 - T_1) F_{T_2}}, \quad \frac{\partial cy}{\partial F_{T_1}} = \frac{1}{(T_2 - T_1) F_{T_1}}, \quad \frac{\partial cy}{\partial r} = 1.$$

Substituting these gives:

$$\Delta cy_{t,(T_1,T_2)} \approx -\frac{1}{(T_2 - T_1) F_{t-1,T_2}} \Delta F_{T_2} + \frac{1}{(T_2 - T_1) F_{t-1,T_1}} \Delta F_{T_1} + \Delta r.$$

This decomposition shows that changes in the forward convenience yield are driven by relative movements in prompt and front futures prices, adjusted for their levels, normalized by maturity, and net of interest rate changes. This structure filters out common trends in the futures curve and reflects regulatory outlook changes, offering a cleaner proxy for carbon transition risk than raw price movements.

Finally, to accounting for the magnitude, we further scale the change in the forward convenience yield by the level of previous convenience yield, which gives the key variable of interest in this research, the return of convenience yield, $R_{cy,t,(T_1,T_2)}$:

$$R_{cy,t,(T_1,T_2)} = \frac{\Delta cy_{t,(T_1,T_2)}}{cy_{t-1,(T_1,T_2)}}.$$

4.3 Portfolio Sorts on Sensitivity to Return in Carbon Convenience Yield

Building on the preceding interpretation, we test the hypothesis that carbon transition risk is a priced factor in European equity markets, using the return of the one-year forward convenience yield as our market-based proxy. An increase in the risk, such as in anticipation of tighter EU ETS regulation, corresponds to a steepening of the futures curve and a more negative forward convenience yield. Consequently, firms with significant exposure to carbon-intensive activities face higher expected compliance costs. In a standard asset pricing framework, if this transition risk is systematic and priced by investors, affected firms should command higher expected returns as compensation for bearing this additional risk.

Our empirical approach employs a two-step portfolio sorting methodology. First, we estimate firm-level factor loadings, or betas ($\beta_i^{r_{cy}}$), by regressing monthly excess stock returns on the monthly return of the one-year forward convenience yield. The time-series regression for each firm i is specified as:

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i^{r_{cy}} \times R_{cy,t} + \epsilon_{i,t},$$

where $R_{i,t}$ is the return of stock i in month t; $R_{f,t}$ is the risk-free rate; and $R_{cy,t}$ is the monthly return of the one-year forward convenience yield in month t. The resulting beta, $\beta_i^{r_{cy}}$, quantifies each firm's sensitivity to our carbon risk proxy.

In the second step, we use the estimated betas to sort stocks into five quintiles. We then calculate the average monthly excess returns for these portfolios, using both value- and equal-weighted schemes. The key variable is the return spread between the highest- and lowest-beta quintiles (High-Low), which we test for statistical significance. To ascertain if this return spread represents compensation for risk beyond that captured by established factors, we regress the portfolio returns on a multi-factor model to estimate alphas:

$$\begin{aligned} R_{j,t} - R_{f,t} = \alpha_j + \sum_{k=1}^{5} \beta_j^{FF_k} \times FF_{k,t} + \beta_j^{MOM} \times MOM_t \\ + \beta_j^{R_{CarbonSpot}} \times R_{CarbonSpot,t} + \beta_j^{R_{OilFront}} \times R_{OilFront,t} + \epsilon_{j,t}, \end{aligned}$$

where $R_{j,t}$ is the return of portfolio j; the $FF_{k,t}$ terms are the five Fama-French European factors (market, size, value, profitability, and investment); MOM_t is the momentum factor; and $R_{CarbonSpot,t}$ and $R_{OilFront,t}$ are the returns on the spot carbon price and front-month oil price, respectively, as .

For our main analysis, we define a 3-year (36-month) rolling lookback period. Betas are estimated using data from the first 35 months of this period, which excludes the month immediately preceding the holding period to prevent any look-ahead bias. Given that our sample spans only 12 years, this estimation period seeks to balance statistical power with the ability to capture time-variation in the betas. The resulting portfolios are held for one month and rebalanced monthly. To ensure our findings are robust to these specifications, we also test alternative configurations, including a 2-year (24-month) estimation period with a one-month holding period and a 3-year (36-month) estimation with a three-month holding period.

Table 1 reports the average monthly excess returns for European STOXX Europ 600 stock portfolios sorted by their sensitivity to the convenience yield return. The central finding is that the high-minus-low spread portfolio yields a consistently negative and statistically significant (at least at 15% level) return. This result holds across the various estimation periods, holding periods, and weighting schemes. Under our primary specification, which uses a 36-month estimation period with sector-neutral, value-weighted portfolios (Panel A, column 3), the high-minus-low spread is -0.51% per month. Furthermore, the quintile portfolios exhibit a general monotonic pattern, with average returns decreasing as beta sensitivity increases. As established, a high beta signifies a hedge against carbon risk, while a low beta signifies exposure to carbon risk. Because the high-minus-low spread is significantly negative, it follows that the portfolio of stocks with higher exposure to carbon transition risk (low beta) is expected to earn a significantly higher return than the portfolio of hedge stocks (high beta). This outperformance is consistent with a carbon risk premium, where investors demand higher compensation for bearing such exposure.

Table 1. Average Monthly Excess Returns of STOXX Europe 600 Stock Portfolios Sorted by Carbon Convenience Yield Return Beta. This table reports average monthly excess returns (%) and t-statistics (in brackets) for quintile portfolios of STOXX Europe 600 stocks. Portfolios are formed monthly by sorting stocks on their estimated beta to the return of the one-year forward carbon convenience yield. The table shows results for both value-weighted and equal-weighted schemes. "General Quintiles" are formed across the entire sample, while "Sector Quintiles" are formed within each industry. Panels A, B, and C show results for different beta estimation and portfolio holding period specifications.

	General Quintiles		*Sector Quintiles*	
Portfolio	Value-weighted	Equal-weighted	Value-weighted	Equal-weighted
Panel A: Three Years Estimation Period with One Month Holding Period				
Low	1.66	1.39	1.52	1.27
	[3.13]	[2.81]	[3.65]	[2.85]
2	1.30	1.20	1.25	1.16
	[3.52]	[3.10]	[3.39]	[3.13]
3	1.19	0.93	1.13	0.98
	[3.49]	[2.66]	[3.53]	[2.78]
4	1.06	0.90	0.90	0.84
	[3.30]	[2.66]	[2.88]	[2.43]
High	0.74	0.69	1.02	0.84
	[2.70]	[2.21]	[3.45]	[2.55]
High–Low	-0.93	-0.69	-0.51	-0.43
	[-1.96]	[-1.98]	[-1.81]	[-1.76]
Panel B: Two Years Estimation Period with One Month Holding Period				
Low	1.63	1.43	1.57	1.36
	[3.47]	[3.13]	[3.97]	[3.22]
2	1.19	1.08	1.17	1.07
	[3.10]	[2.84]	[3.38]	[3.07]
3	1.19	1.03	1.04	0.97
	[3.79]	[3.21]	[3.49]	[2.92]
4	1.08	0.98	1.05	0.96
	[3.60]	[3.09]	[3.41]	[3.03]
High	0.85	0.86	0.97	1.00
	[2.94]	[2.67]	[3.32]	[3.04]
High–Low	-0.78	-0.57	-0.60	-0.35
	[-1.83]	[-1.63]	[-2.51]	[-1.51]
Panel C: Three Years Estimation Period with Three Month Holding Period				
Low	1.68	1.41	1.50	1.29
	[3.15]	[2.90]	[3.65]	[2.92]
2	1.34	1.20	1.27	1.16
	[3.56]	[3.09]	[3.46]	[3.21]
3	1.07	0.87	1.06	0.97
	[3.00]	[2.47]	[3.15]	[2.65]
4	1.06	0.91	0.93	0.83
	[3.49]	[2.67]	[2.98]	[2.35]
High	0.75	0.71	0.99	0.84
	[2.69]	[2.26]	[3.40]	[2.61]
High–Low	-0.93	-0.71	-0.51	-0.46
	[-1.95]	[-2.11]	[-1.79]	[-1.87]

We next examine whether this return premium persists after controlling for exposures to known risk factors. Table 2 presents the resulting alphas and factor loadings from the multi-factor regression for our primary specification. The key result, shown in the last column of the table, is that the high-minus-low portfolio generates a statistically significant alpha of -0.77% per month. This demonstrates that the outperformance of high carbon risk exposure stocks is not explained by exposure to standard market, size, value, profitability, or investment factors, nor by momentum or direct energy price shocks. The persistence of this alpha suggests that our convenience yield return proxy captures a priced risk dimension distinct from other documented factors in the European equity market.

Table 2. Multi-Factor Model Regressions for STOXX Europe 600 Stock Portfolios Sorted by Carbon Convenience Yield Return Beta. This table reports the monthly alphas (%) and factor loadings from multi-factor time-series regressions of the quintile portfolio excess returns. The portfolios are the value-weighted, sector-neutral quintiles from our main specification (3-year estimation period and 1-month holding period), which are formed by sorting stocks on their Carbon Convenience Yield Beta. The independent variables in the regression model include the Fama-French Europe five factors (MKT, SMB, HML, RMW, CMA), a momentum factor (MOM), and the returns on spot carbon and front-month oil prices. Reported values are the regression coefficients and their corresponding t-statistics (in brackets).

	Low	2	3	4	High	High-Low
α	1.06	0.88	0.58	0.41	0.45	-0.77
	[5.25]	[5.36]	[5.05]	[1.89]	[2.47]	[-2.86]
MKT	0.79	0.69	0.66	0.59	0.56	-0.24
	[15.30]	[13.88]	[15.87]	[13.79]	[13.89]	[-3.71]
SMB	-0.04	-0.17	-0.34	-0.27	-0.16	-0.10
	[-0.37]	[-1.58]	[-3.32]	[-2.44]	[-1.29]	[-0.66]
HML	0.36	0.03	0.16	0.03	0.07	-0.29
	[2.66]	[0.22]	[1.35]	[0.28]	[0.71]	[-1.96]
RMW	0.10	0.05	0.30	0.36	0.31	0.24
	[0.51]	[0.29]	[2.12]	[2.05]	[1.58]	[0.78]
CMA	-0.14	0.06	-0.05	-0.09	-0.12	0.02
	[-0.57]	[0.30]	[-0.37]	[-0.56]	[-0.65]	[0.06]
MOM	-0.23	-0.21	-0.07	-0.14	-0.08	0.15
	[-2.69]	[-2.72]	[-0.95]	[-2.19]	[-1.60]	[1.63]
$\Delta CarbonSpot$	-0.02	-0.02	0.00	0.02	0.04	0.06
	[-1.46]	[-1.35]	[0.26]	[1.27]	[2.47]	[3.16]
$\Delta OilFront$	0.02	0.03	0.03	0.03	0.03	0.01
	[1.20]	[2.01]	[2.76]	[1.83]	[1.72]	[0.70]

Having established a significant risk premium in the broad market, we next test whether this pricing is driven by the firms most directly impacted by the regulation. This analysis reveals a divergence: the high-minus-low portfolio for the ETS-exposed subsample yields both a mean return and a multi-factor alpha that are statistically insignificant (Table 3). Hence, our market-implied proxy is not priced in the subsample of firms with direct EU ETS exposure.

Table 3. Returns and Alphas for EU ETS-Exposed STOXX Europe 600 Stock Portfolios Sorted by Carbon Convenience Yield Return Beta. This table repeats the portfolio analysis for the subsample of STOXX Europe 600 firms with direct operational exposure to the EU Emissions Trading System (EU ETS). It reports results for value-weighted, sector-neutral quintile portfolios sorted on their Carbon Convenience Yield Beta. The table presents both the average monthly excess returns and multi-factor alphas. The alphas and corresponding factor loadings are estimated from a regression model that includes the Fama-French Europe five factors (MKT, SMB, HML, RMW, CMA), a momentum factor (MOM), and the returns on spot carbon and front-month oil prices. All reported values are coefficients with their corresponding t-statistics shown in brackets.

	Low	2	3	4	High	High-Low
$R_{j,t} - R_{f,t}$	0.98	1.09	0.95	0.75	1.10	0.12
	[2.71]	[2.62]	[3.37]	[2.21]	[4.49]	[0.66]
α	0.38	0.61	0.33	0.20	0.46	-0.08
	[1.85]	[2.59]	[1.62]	[1.10]	[2.44]	[-0.38]
MKT	0.77	0.70	0.55	0.54	0.47	-0.31
	[11.24]	[12.31]	[11.80]	[10.30]	[10.32]	[-5.35]
SMB	-0.15	-0.21	-0.53	-0.28	-0.21	-0.03
	[-1.51]	[-1.48]	[-5.04]	[-2.26]	[-1.30]	[-0.19]
HML	0.49	0.30	0.02	0.18	0.32	-0.18
	[3.07]	[2.03]	[0.10]	[1.30]	[2.59]	[-1.37]
RMW	0.23	0.39	0.47	0.50	0.77	0.56
	[1.18]	[1.85]	[2.17]	[2.46]	[5.05]	[3.30]
CMA	0.11	0.28	0.61	0.23	0.15	0.04
	[0.46]	[1.02]	[2.95]	[1.20]	[0.80]	[0.20]
MOM	-0.12	-0.18	0.03	-0.18	-0.04	0.08
	[-1.51]	[-2.03]	[0.53]	[-1.59]	[-0.45]	[1.05]
$\Delta CarbonSpot$	0.00	-0.02	0.02	0.02	0.04	0.04
	[0.17]	[-1.30]	[1.44]	[1.39]	[1.91]	[1.84]
$\Delta OilFront$	0.01	0.06	0.05	0.07	0.04	0.04
	[0.44]	[2.25]	[1.89]	[4.23]	[2.08]	[1.69]

A potential statistical explanation for this insignificant result is that these firms, being uniformly exposed to carbon regulation, might exhibit less cross-sectional dispersion in their risk sensitivities, diminishing the power of our portfolio sorts. To examine this, we analyze the distribution of the estimated betas. As shown in Table 4, this hypothesis is not supported by the data. Despite the smaller sample size, the standard deviation of the betas for the ETS subsample is nearly identical to that of the full STOXX Europe 600 sample. Therefore, a lack of beta dispersion can be ruled out as a primary explanation for the absence of a premium.

Table 4. Distribution of Sample Carbon Convenience Yield Return Beta. This table reports the distribution of betas with 3-year estimation and 1-month holding period rolling window, for both full sample of STOXX Europe 600 stocks and the subsample with EU ETS exposure.

	STOXX Europe 600	STOXX Europe 600 with EU ETS Exposure
Mean	-0.059	-0.055
SD	0.094	0.094
Min	-0.606	-0.586
p25	-0.111	-0.104
Median	-0.051	-0.048
p75	0.002	0.009
Max	0.461	0.324
N	59,890	13,505

With a simple statistical explanation ruled out, this divergence suggests that what is broadly termed as carbon risk manifests in economically distinct ways across the market. For the broad STOXX Europe 600, the premium likely reflects the pricing of a systematic, long-term transition risk, as investors reward potential beneficiaries and penalize firms with uncertain future climate liabilities. For the ETS-regulated firms, however, this risk appears to be treated not as a priced systematic factor but as a direct operational cost. This phenomenon may occur if firms are able to pass compliance costs on to customers, absorb them through operational adjustments, or if investors, while aware of the costs, do not view them as a source of undiversifiable risk. This distinction motivates the next step of our analysis: shifting from market-based expectations to realized accounting burdens, and asking whether the actual financial impact of EU ETS compliance is reflected in expected stock returns.

5 Corporate Carbon Expense

5.1 Ratios of Carbon Expense

The EU ETS requires installations to surrender a number of allowances equal to their verified annual emissions by the end of April of the following year. These compliance transactions are recorded in the Union Registry. Because transaction-level prices are not publicly available, we proxy the effective carbon price using the annual average of EUA daily futures prices. The firm-level carbon expense is computed by multiplying the number of surrendered allowances by this annual average price. Free allocations are not netted out, since firms that emit below their cap can sell the surplus, making the full surrender volume an appropriate measure of economic cost.

To assess the materiality of carbon expenses, we scale each firm's estimated annual carbon expense by net revenue, gross income, and operating income. These three benchmarks reflect different layers of the income statement and capture increasing levels of profitability. Although carbon compliance costs are not separately reported, they may be embedded in various line items, depending on how firms account for them. The stylized income statement below illustrates how carbon expenses could be treated depending on their operational or administrative classification:

	Gross Sales
–	Returns, Discounts, and Allowances
=	**Net Revenue (Net Sales)**
–	Cost of Goods Sold (COGS)
	(Carbon allowance cost may be included here if production-related)
–	Depreciation/Depletion/Amortization (production-related)
=	**Gross Income**
–	Selling, General & Administrative Expenses (SG&A)
	(Carbon allowance cost may be included here if treated as compliance/admin)
–	R&D and Other Operating Expenses
–	Depreciation/Amortization (non-production)
=	**Operating Income**

This mapping highlights that the financial burden of carbon compliance can appear in different places depending on whether it is treated as a production input or an administrative cost. By evaluating expense ratios relative to multiple profitability measures, we capture a broader view of firms' carbon expense intensity.

One caveat in interpreting these ratios is that carbon expenses are constructed using emissions data recorded at the installation (plant) level and then aggregated to the firm level. However, without detailed information on each company's internal structure and financial reporting boundaries, it is not possible to consolidate these costs in the same manner as official financial statements. The analysis therefore assumes that all operator accounts linked to a firm fully contribute their carbon expenses to that firm. While this approach may overstate or understate true consolidated costs for certain firms, it provides a consistent and replicable method for estimating firm-level carbon exposure across the sample.

Constructing a firm-level measure of carbon expense from public data presents two primary methodological challenges: the aggregation of emissions and the pricing of surrendered allowances. We address these challenges with a focus on creating the most transparent and replicable proxy possible. First, as the Union Registry links emissions to operator accounts and not directly to the publicly listed parent company , our approach uses a rigorous two-step matching procedure to map these entities. While this necessary aggregation may not perfectly mirror a firm's complex internal accounting boundaries, it provides a consistent framework for estimation. Second, since transaction-level prices are not disclosed, we proxy the cost using the annual average of daily futures prices to create a standardized measure of the cost burden. Despite these necessary assumptions, this "bottom-up" methodology, grounded in audited operational data, offers a distinct advantage over commonly used third-party ESG ratings, which often rely on proprietary estimation models and have been criticized for their opacity and potential for bias. Therefore, while our measure is subject to the limitations of public data availability, we contend that it provides a transparent and robust estimate of realized carbon costs.

Figure 6 presents the distribution of carbon expense ratio to operating income over the sample period. The figure includes two panels: the upper panels show annual trends in the cross-sectional distribution, while the lower panels break down the ratios by sector. Four sectors, Consumer Discretionary, Financials, Information Technology, and Communication Services, are omitted due to limited data and minimal relevance in terms of carbon expense exposure. The figures for carbon expense ratios to net revenue and gross income exhibit similar trends and are omitted for brevity.

All three ratios exhibit a rising trend over time, broadly mirroring the increase in EUA prices documented in Figure 4. By 2023, firms at the 75th percentile of the distribution faced carbon expenses equivalent to 1.41% of net revenue, 5.10% of gross income, and 12.42% of operating income. These magnitudes imply that, in the absence of the EU ETS, a quarter of the sample firms could see gross income increase by over 5% and operating income by more than 12%. These are not marginal effects, and they indicate that carbon compliance imposes a material burden on profitability.

The upward trend and elevated ratios are especially pronounced in the Materials and Utilities sectors. In these sectors, more than half of the sampled firms have carbon expenses exceeding 1% of revenue, 5% of gross income, and 10% of operating income. These findings suggest that carbon expenses have a significant and persistent impact on firms' operating margins, particularly in industries that are directly targeted by emissions regulation. This underscores the economic relevance of carbon risk, even if it does not yet appear to be priced in stock returns.

5.2 *Portfolios Sorts on Ratios of Carbon Expense*

Having shown that carbon expenses represent a nontrivial burden on firm performance, we now ask whether this burden is priced in expected stock returns. To do so, we conduct a standard portfolio sorting exercise. Each year, firms with EU ETS operator accounts are sorted into quintiles based on their carbon expense intensity (using each of the three

Figure 6. Ratio of Carbon Expense to Operating Income. This figure shows the annual distribution of the ratio of carbon expense to operating income across firms. Whiskers represent the 10th and 90th percentiles.

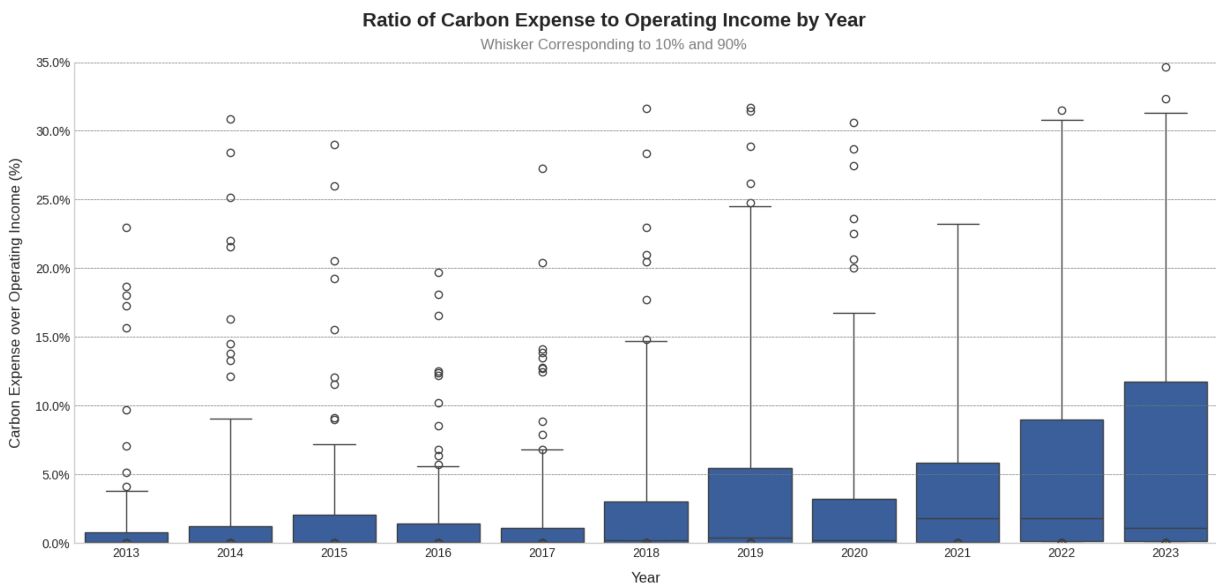

(a) Overall Distribution.

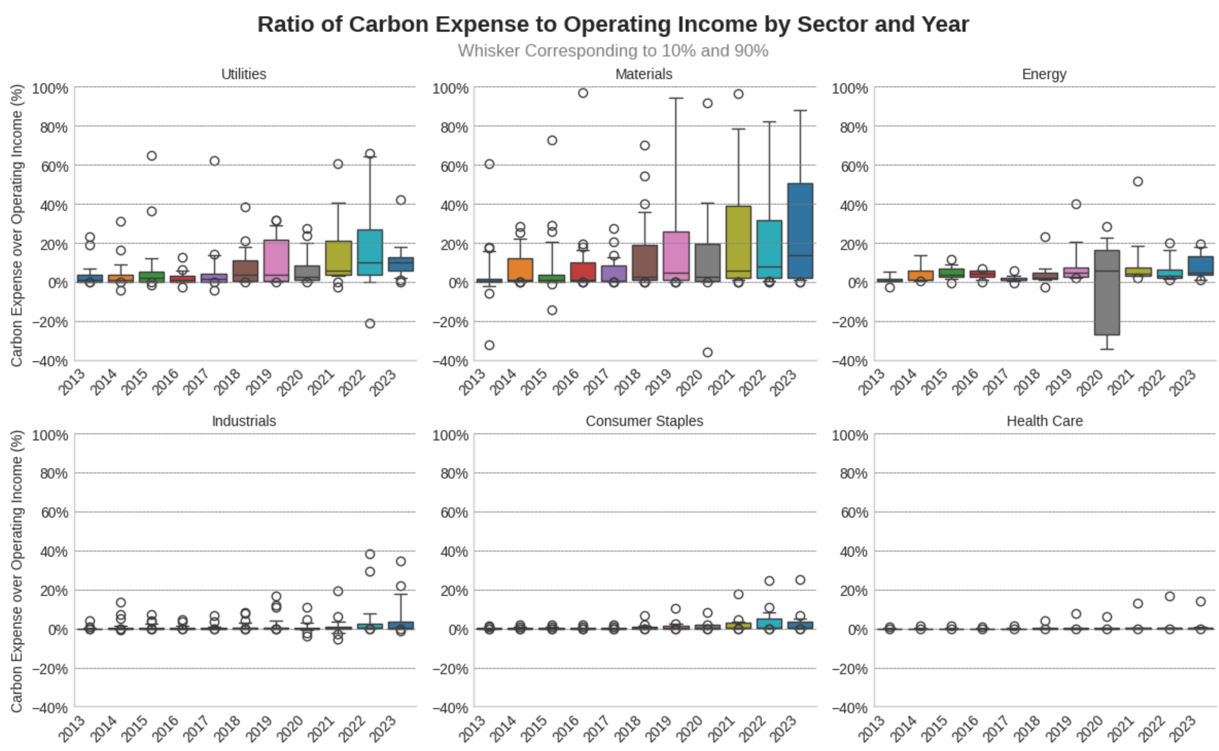

(b) Distribution by Sector. Sectors are ordered by mean ratio in the final year. Four sectors (Consumer Discretionary, Financials, Information Technology, and Communication Services) are excluded due to data limitations and lack of significance.

expense ratios). Portfolios are rebalanced annually, with sorting based on lagged financial and emissions data in line with the EU ETS compliance cycle. This design ensures portfolio formation uses only information that would have been available to investors at the time.

Portfolios are rebalanced at the end of April each year. The carbon expense ratios used for sorting are based on emissions and financial data from the previous calendar year, which would not be available before the end of March. This timing aligns with EU ETS reporting requirements, as installations must submit their audited emissions data for year t by the end of March in year $t+1$. As a result, quintile portfolios are formed using lagged information and held from May of year $t+1$ through April of year $t+2$.[6]

With the portfolios constructed, average returns for each quintile are calculated using both value-weighted and equal-weighted schemes. The return spread between the highest and lowest quintiles is then tested for statistical significance. To further examine whether these return differences reflect compensation for risk, we regress portfolio returns on the five European Fama-French Europe factors to estimate factor-adjusted alphas and exposures.

To control for industry composition, we also conduct sector-neutral sorts by forming quintiles within each sector before aggregating portfolios. This adjustment prevents carbon-intensive sectors from dominating the highest-cost portfolios.

Table 5 reports average excess returns for value-weighted and equal-weighted portfolios. Across all specifications, we find no consistent relationship between carbon expense intensity and returns. High-minus-low portfolio spreads are statistically insignificant, and there is no monotonic pattern across quintiles. These results hold for both general and sector-adjusted sorts and across all three expense ratios.

Table 6 presents alphas and factor loadings from Fama-French five-factor regressions for the value-weighted, sector-neutral portfolios, corresponding to column (3) in Table 5. Consistent with the excess return results, none of the alphas from the high-minus-low portfolios are statistically significant. However, many individual quintile portfolios, particularly in

[6] For instance, emissions and financial data from 2013 are used to sort stocks into portfolios in April 2014, which are then held from May 2014 to April 2015.

Table 5. Average Excess Returns Sorted by Carbon Expense Ratios. This table reports the average excess returns and corresponding t-statistics (in square brackets) for portfolio quintiles sorted by various carbon intensity measures. Columns (1) and (2) show results from general quintile sorts, while Columns (3) and (4) report results from within-sector quintile sorts. Columns (1) and (3) are value-weighted portfolios, and Columns (2) and (4) are equal-weighted portfolios.

	General Quintiles		*Sector Quintiles*	
Portfolio	Value-weighted	Equall-weighted	Value-weighted	Equall-weighted
Panel A: Ratio of Carbon Expense to Net Revenue				
Low	0.98	0.78	0.71	0.68
	[3.20]	[2.09]	[2.36]	[2.00]
2	0.90	0.69	1.05	0.70
	[3.19]	[1.89]	[4.12]	[2.16]
3	1.01	0.93	0.87	0.78
	[3.23]	[2.74]	[2.83]	[2.10]
4	0.76	0.59	0.95	0.86
	[1.98]	[1.69]	[2.76]	[2.25]
High	1.21	1.00	1.09	0.92
	[3.53]	[2.38]	[2.81]	[2.22]
High–Low	0.23	0.22	0.38	0.24
	[0.97]	[1.02]	[1.59]	[1.20]
Panel B: Ratio of Carbon Expense to Gross Income				
Low	0.87	0.71	0.89	0.68
	[3.26]	[2.28]	[3.40]	[2.19]
2	1.03	0.80	0.81	0.64
	[3.23]	[2.30]	[2.92]	[2.03]
3	1.02	0.77	1.08	0.92
	[2.44]	[1.93]	[3.69]	[2.69]
4	0.68	0.76	0.82	0.85
	[1.81]	[2.15]	[2.16]	[2.01]
High	1.32	0.96	1.20	0.90
	[3.63]	[2.25]	[3.13]	[2.21]
High–Low	0.44	0.25	0.31	0.23
	[1.37]	[0.86]	[1.16]	[1.16]
Panel C: Ratio of Carbon Expense to Operating Income				
Low	1.08	0.77	0.87	0.64
	[3.33]	[1.96]	[2.98]	[1.85]
2	0.75	0.77	0.77	0.72
	[2.84]	[2.28]	[2.98]	[2.33]
3	1.10	0.82	0.99	0.87
	[3.45]	[2.49]	[3.13]	[2.40]
4	0.68	0.81	0.91	0.80
	[1.78]	[2.20]	[2.61]	[2.06]
High	1.16	0.85	1.26	0.92
	[3.30]	[1.95]	[3.21]	[2.16]
High–Low	0.08	0.08	0.38	0.28
	[0.31]	[0.40]	[1.44]	[1.55]

the upper and middle ranges, exhibit significantly positive excess returns and corresponding alphas. This suggests that firms in the STOXX Europe 600 index with EU ETS operator accounts may earn higher average returns than predicted by the Fama-French Europe five-factor model. Nevertheless, despite the non-negligible burden of carbon expenses relative to revenue and income, the carbon expense ratios themselves do not explain variation in expected returns across firms. Results for alternative sorting and weighting schemes yield similarly insignificant pricing evidence and are omitted for brevity.

As a side observation, we note modest negative loadings on the investment (CMA) factor and weak positive loadings on the value (HML) factor in the long-short portfolio. Although not the focus of this analysis, these patterns may reflect underlying firm characteristics. For instance, firms with higher carbon expenses may rely more heavily on tangible capital in production, which may correlate with higher book-to-market ratios and emissions intensity. Similarly, companies whose profits are more heavily eroded by carbon costs may be incentivised to invest more aggressively in green technologies, in line with the objectives of the EU ETS.

5.3 Double Sorting on Carbon Expense Ratios then Carbon Convenience Yield Return Sensitivity

The preceding analyses establish that for firms directly exposed to the EU ETS, neither our market-based measure of systematic risk nor our accounting-based measure of financial burden is priced in the cross-section of returns. This leads us to test an interaction hypothesis: a risk premium for systematic carbon risk may exist, but only for the group of firms where the financial impact of carbon compliance is most severe. The economic rationale is that investors may only price a firm's exposure to market-wide carbon risk when its firm-specific, realized costs are high enough to be financially material. For firms with negligible carbon expense ratios, this systematic risk may be less relevant, whereas for firms facing a high cost burden, it should command a premium.

Table 6. Alphas and Factor Loadings from Five-Factor Regressions Sorted by Carbon Expense Ratios with Sector Neutral Value Weighted Portfolios. This table reports portfolio alphas and Fama-French Europe five-factor loadings from time-series regressions. Each panel presents results for quintile portfolios sorted by different carbon intensity measures. Reported values are coefficients and corresponding t-statistics (in square brackets).

	α	MKT	SMB	HML	RMW	CMA
Panel A: Ratio of Carbon Expense to Net Revenue						
Low	0.29	0.59	-0.39	0.22	0.78	0.35
	[1.70]	[10.77]	[-3.75]	[1.47]	[3.78]	[1.62]
2	0.66	0.60	-0.61	0.34	0.59	0.21
	[3.46]	[11.30]	[-5.24]	[2.08]	[2.84]	[0.71]
3	0.41	0.66	-0.33	0.61	0.98	0.38
	[2.14]	[11.13]	[-1.85]	[3.73]	[4.42]	[1.52]
4	0.46	0.82	-0.26	0.59	0.64	0.05
	[2.15]	[15.82]	[-1.37]	[5.13]	[4.28]	[0.22]
High	0.60	0.90	-0.34	0.38	0.36	-0.23
	[3.86]	[21.25]	[-3.03]	[2.82]	[1.83]	[-0.89]
High–Low	0.17	0.31	0.07	0.17	-0.39	-0.58
	[0.87]	[5.20]	[0.55]	[1.02]	[-1.57]	[-2.35]
Panel B: Ratio of Carbon Expense to Gross Income						
Low	0.50	0.56	-0.50	0.07	0.65	0.34
	[2.76]	[12.68]	[-4.48]	[0.51]	[3.12]	[1.50]
2	0.40	0.57	-0.60	0.26	0.68	0.24
	[2.18]	[10.36]	[-4.70]	[1.47]	[3.21]	[0.73]
3	0.63	0.62	-0.27	0.68	0.99	0.25
	[3.27]	[9.19]	[-1.80]	[4.29]	[5.53]	[1.11]
4	0.31	0.87	-0.35	0.64	0.76	0.09
	[1.43]	[15.72]	[-1.88]	[3.87]	[3.65]	[0.39]
High	0.72	0.87	-0.35	0.45	0.30	-0.31
	[3.57]	[16.13]	[-2.25]	[3.25]	[1.69]	[-1.27]
High–Low	0.08	0.30	0.17	0.39	0.31	-0.64
	[0.38]	[5.52]	[1.14]	[2.18]	[-1.37]	[-2.76]
Panel C: Ratio of Carbon Expense to Operating Income						
Low	0.46	0.61	-0.44	0.18	0.71	0.34
	[2.37]	[11.23]	[-3.52]	[1.14]	[3.32]	[1.60]
2	0.36	0.55	-0.50	0.34	0.76	0.20
	[2.15]	[9.98]	[-4.88]	[2.10]	[3.98]	[0.78]
3	0.53	0.71	-0.32	0.51	0.88	0.39
	[2.95]	[10.84]	[-1.92]	[2.79]	[3.91]	[1.24]
4	0.42	0.81	-0.56	0.62	0.59	-0.12
	[2.13]	[14.07]	[-3.19]	[3.89]	[3.05]	[-0.52]
High	0.80	0.83	-0.22	0.57	0.45	-0.06
	[3.76]	[15.56]	[-1.35]	[3.04]	[2.16]	[-0.18]
High–Low	0.20	0.22	0.24	0.40	-0.23	-0.40
	[0.87]	[2.76]	[1.51]	[1.83]	[-0.81]	[-1.30]

To test this interaction hypothesis, we employ a 2×2 double sorting method, a choice informed by the limited sample size of 129 firms with EU ETS exposure. The procedure first sorts firms into high and low groups based on their carbon expense ratio to operating income. Subsequently, within each of these initial groups, firms are sorted again into high and low beta portfolios based on their sensitivity to the return in carbon convenience yield. The technical details of this portfolio formation are consistent with the main specifications used in Sections 4.3 and 5.2, including the use of a 3-year estimation period for beta, sector-neutral, value-weighted portfolios, and a one-month holding period.

The results of the double sort are presented in Table 7. The mean returns and multi-factor alphas for the high-minus-low beta spread portfolios are statistically insignificant for both the high and low expense ratio groups. Therefore, we reject the hypothesis that a risk premium for systematic carbon risk is conditionally priced based on the level of a firm's carbon expense ratio.

Table 7. Returns and Alphas for EU ETS-Exposed STOXX Europe 600 Stock Portfolios Double Sorted by Ratio of Carbon Expense to Operating Income then Carbon Convenience Yield Return Beta. This table reports the portfolio analysis for the subsample of STOXX Europe 600 firms with direct operational exposure to the EU Emissions Trading System (EU ETS). It reports results for value-weighted, sector-neutral quintile portfolios first sorted on carbon expense ratio to operating income, then on carbon convenience yield return beta. The table presents both the average monthly excess returns and multi-factor alphas. The alphas and corresponding factor loadings are estimated from a regression model that includes the Fama-French Europe five factors (MKT, SMB, HML, RMW, CMA), a momentum factor (MOM), and the returns on spot carbon and front-month oil prices. All reported values are coefficients with their corresponding t-statistics shown in brackets.

	High Expense Spread	Low Expense Spread
$R_{j,t} - R_{f,t}$	0.15	-0.16
	[0.52]	[-0.76]
α	-0.20	-0.23
	[-0.76]	[-1.21]

Therefore, the results from both the single and double sorting procedures consistently show that for the firms most directly impacted by the EU ETS, neither our market-based nor our accounting-based measures of carbon risk are priced in the cross-section of returns. This reinforces the finding of a disconnect between the operational impact of carbon regulation and its reflection in equity valuations for this key group of firms.

6 Conclusion

This study examines whether carbon risk is priced in the European equity market using two complementary and transparent measures. The first is a forward-looking, market-based proxy derived from the one-year forward convenience yield in the carbon futures market, which reflects investor expectations about future allowance prices. The second is an accounting-based proxy, the carbon expense ratio, constructed by aggregating firm-level verified emissions from the EU's Union Registry and multiplying them by the average annual price of carbon allowances. We find that both proxies capture economically meaningful risks. Our accounting-based measure reveals a significant and growing financial burden, with carbon expenses for the most-exposed quartile of firms reaching 12.4% of operating income by 2023. These are not marginal effects; they represent a material drag on profitability. Concurrently, our market-based proxy provides a clear, forward-looking signal of expected regulatory and compliance costs.

A central finding of our analysis is the disconnect between the clear materiality of these risks and their reflection in equity valuations for the very firms most affected. Although sensitivity to our market-based proxy is priced in the broad STOXX Europe 600, generating a significant monthly alpha of -0.77% for a high-minus-low portfolio, this premium vanishes within the subsample of firms directly regulated by the EU ETS. Furthermore, the realized impact of their carbon expense intensity appears to be neglected as well, showing no association with a return premium across multiple specifications.

Given that firm-level explanations such as cost pass-through or sector-specificity appear insufficient to resolve this finding, a compelling avenue for future research is to investigate this apparent investor neglect. An important open question is why investors seem to disregard these material, verifiable costs for ETS-regulated firms. Future work could test for specific mechanisms, such as investor inattention, the perception that these risks are idiosyncratic rather than systematic, or a form of market segmentation where the specialized investor base for these industrial firms focuses on other operational drivers.

Overall, these findings highlight a disconnect between the operational impact of carbon costs and their recognition in equity pricing. As regulatory stringency increases and investor attention to climate-related risks grows, the framework developed in this study provides a transparent and replicable basis for monitoring whether carbon risk eventually becomes reflected in expected stock returns. The absence of a carbon premium for EU ETS-exposed firms, despite substantial and rising cost burdens, underscores a gap between regulatory exposure and investor pricing, a disconnect with direct implications for ESG integration and the responsiveness of equity markets to climate policy.

References

Afonin, A., Bredin, D., Cuthbertson, K., Muckley, C., Nitzsche, D., 2018. Carbon Portfolio Management. International Journal of Finance & Economics 23, 349–361.

Akey, P., Appel, I., 2021. The Limits of Limited Liability: Evidence from Industrial Pollution. The Journal of Finance 76, 5–55.

Alberola, E., Chevallier, J., Chèze, B., 2008. Price Drivers and Structural Breaks in European Carbon Prices 2005–2007. Energy Policy 36, 787–797.

Ardia, D., Bluteau, K., Boudt, K., Inghelbrecht, K., 2023. Climate Change Concerns and the Performance of Green vs. Brown Stocks. Management Science 69, 7607–7632.

Aswani, J., Raghunandan, A., Rajgopal, S., 2024. Are Carbon Emissions Associated with Stock Returns? Review of Finance 28, 75–106.

Azlen, M., Gostlow, G., Child, A., 2022. The Carbon Risk Premium. The Journal of Alternative Investments 25, 33–64.

Batten, J. A., Maddox, G. E., Young, M. R., 2021. Does Weather, or Energy Prices, Affect Carbon Prices? Energy Economics 96, 105016.

Berg, F., Fabisik, K., Sautner, Z., 2021. Is History Repeating Itself? The (Un)Predictable Past of ESG Ratings. Working Paper.

Berg, F., Kölbel, J. F., Rigobon, R., 2022. Aggregate Confusion: The Divergence of ESG Ratings. Review of Finance 26, 1315–1344.

Bolton, P., Kacperczyk, M., 2021. Do Investors Care about Carbon Risk? Journal of Financial Economics 142, 517–549.

Bredin, D., Parsons, J., 2016. Why is Spot Carbon so Cheap and Future Carbon so Dear? The Term Structure of Carbon Prices. The Energy Journal 37, 83–107.

Chevallier, J., 2009. Carbon Futures and Macroeconomic Risk Factors: A View from the EU ETS. Energy Economics 31, 614–625.

Demiralay, S., Gencer, H. G., Bayraci, S., 2022. Carbon Credit Futures as an Emerging Asset: Hedging, Diversification and Downside Risks. Energy Economics 113, 106196.

Engle, R. F., Giglio, S., Kelly, B., Lee, H., Stroebel, J., 2020. Hedging Climate Change News. The Review of Financial Studies 33, 1184–1216.

Eskildsen, M., Ibert, M., Jensen, T. I., Pedersen, L. H., 2024. In Search of the True Greenium.

ESMA, 2022. Final Report on Emission Allowances and Associated Derivatives. ESMA 70-445-38.

Fuchs, M., Stroebel, J., Terstegge, J., 2024. Carbon VIX: Carbon Price Uncertainty and Decarbonization Investments. NBER Working Paper .

Galtung, J., Ruge, M. H., 1965. The Structure of Foreign News. Journal of Peace Research 2, 64–91.

Harcup, T., and O'Neill, D., 2001. What Is News? Galtung and Ruge Revisited. Journalism Studies 2, 261–280.

Hsu, P.-H., Li, K., Tsou, C.-Y., 2023. The Pollution Premium. The Journal of Finance 78, 1343–1392.

Lutz, B. J., Pigorsch, U., Rotfuß, W., 2013. Nonlinearity in Cap-and-Trade Systems: The EUA Price and its Fundamentals. Energy Economics 40, 222–232.

Mansanet-Bataller, M., Pardo, A., Valor, E., 2007. CO_2 Prices, Energy and Weather. The Energy Journal 28, 73–92.

Matsumura, E. M., Prakash, R., Vera-Munoz, S. C., 2014. Firm-Value Effects of Carbon Emissions and Carbon Disclosures. The Accounting Review 89, 695–724.

Medina, V., Pardo, A., 2013. Is the EUA a New Asset Class? Quantitative Finance 13, 637–653.

Palao, F., Pardo, Á., 2021. The Inconvenience Yield of Carbon Futures. Energy Economics 101, 105461.

Pedersen, L. H., Fitzgibbons, S., Pomorski, L., 2021. Responsible Investing: The ESG-Efficient Frontier. Journal of Financial Economics 142, 572–597.

Sautner, Z., Van Lent, L., Vilkov, G., Zhang, R., 2023. Firm-Level Climate Change Exposure. The Journal of Finance 78, 1449–1498.

Swinkels, L., Yang, J., 2023. Investing in Carbon Credits. The Journal of Alternative Investments 26, 28–59.

Task Force on Climate-Related Financial Disclosures, 2017. TCFD Recommendations. https://www.fsb-tcfd.org/recommendations/.

Trück, S., Weron, R., 2016. Convenience Yields and Risk Premiums in the EU-ETS—Evidence from the Kyoto Commitment Period. Journal of Futures Markets 36, 587–611.

Ye, J., Xue, M., 2021. Influences of Sentiment from News Articles on EU Carbon Prices. Energy Economics 101, 105393.

Zhang, S., 2025. Carbon Returns across the Globe. The Journal of Finance 80, 615–645.

Zhu, B., Ye, S., Han, D., Wang, P., He, K., Wei, Y.-M., Xie, R., 2019. A Multiscale Analysis for Carbon Price Drivers. Energy Economics 78, 202–216.