Beyond Emissions:

Regulatory Carbon Positions and Stock Returns

Zong-Wei Yeh* Wei-Che Tsai† Austin Liu‡ Jie Luo§

This version: October 2026

Abstract

The allocation shortfall, the share of a firm's verified emissions not covered by its free allowances under the EU Emissions Trading System, measures how tightly the cap binds the firm relative to the allocation the regulator grants. Among European listed firms outside the utility sector, firms whose emissions persistently fall short of their allocation earn higher subsequent returns than firms whose emissions persistently exceed it: a value-weighted spread of 5.8% per year over 2009–2024, with an alpha of 6.1% per year after the Fama–French five factors and momentum. The spread is not explained by carbon-price exposure or carbon costs, and it is concentrated among firms for which the system is economically small. Controlling for firm characteristics, low-shortfall firms subsequently beat analysts' earnings forecasts. The evidence is consistent with investors and analysts not fully incorporating the information in firms' regulatory carbon positions, rather than with compensation for carbon-price risk.

Keywords: EU Emissions Trading System, allocation shortfall, information processing, underreaction, cross-section of stock returns.
JEL codes: G12, G14, Q54, Q58.

*Department of Finance, National Chung Cheng University, Chiayi, Taiwan. Email: zweyeh@ccu.edu.tw.
†Corresponding author. Department of Finance, National Sun Yat-sen University, Kaohsiung, Taiwan. Email: weiche@mail.nsysu.edu.tw.
‡UCLA Anderson School of Management, University of California, Los Angeles, USA. Email: austin.liu.phd@anderson.ucla.edu.
§Department of Finance, National Sun Yat-sen University, Kaohsiung, Taiwan. Email: 134030004@student.nsysu.edu.tw.

1 Introduction

The European Union Emissions Trading System (EU ETS) requires each regulated installation to surrender one allowance for every tonne of carbon dioxide it emits, and it hands out a large part of these allowances for free.[1] How many free allowances a firm receives relative to what it emits defines its regulatory carbon position: how tightly the cap binds the firm, and how much of its emissions it must cover at the market price of carbon. Recent work shows that this position shapes how stock prices respond to the carbon market. Bolton et al. (2026) find that, in Phase 3 of the EU ETS, a rise in the carbon price is associated with lower stock prices for firms with a larger allowance shortfall and with higher stock prices for firms with greater permit coverage. Laurs et al. (2026) find this interaction in every phase after the pilot phase and show that firms with larger allocation shortfalls comply less often, while announcements of non-compliance lower stock prices.[2] In these studies the position matters through its interaction with carbon-price changes or compliance news. In this paper, we ask whether the position itself predicts the cross-section of subsequent stock returns, and if so, why.

Following Laurs et al. (2026), we measure the position with the *allocation shortfall*, the share of a firm's verified emissions that its free allocation does not cover. Both inputs are recorded by the regulator: accredited verifiers audit the emissions, and the allocation is legally binding. The allocation follows rules. Since 2013 it has been computed from EU-wide product benchmarks applied to an installation's historical activity, with more generous treatment for sectors deemed exposed to carbon leakage, and before 2013 it largely

[1] The EU ETS covers more than ten thousand installations in the EU member states, Iceland, Liechtenstein and Norway, which together account for about 40% of greenhouse gas emissions in Europe (Akey et al. 2025). Appendix B describes the trading phases and the allocation rules.

[2] Millischer et al. (2023) relate weekly carbon-price changes to the returns of 338 European firms from 2013 to 2021 and find that firms with higher carbon costs perform worse when carbon prices rise. Laurs et al. (2026) link more than 15,300 stationary installations to more than 1,200 listed firms from 2005 to 2023, use daily returns, and find cumulative abnormal returns of at least −2% over the five days around announcements of non-compliance. Bolton et al. (2026) also find that firms with a significant permit shortfall reduce their verified emissions in the EU.

followed historical emissions.[3] The shortfall therefore combines a firm's operations with the regulatory treatment it receives. Much of its cross-sectional variation is slow-moving, because benchmarks, leakage status and production technology change slowly, while part of it reflects year-to-year fluctuations in output and weather. We therefore separate a persistent component, the firm's slow-moving regulatory position, from a transitory component.

Economic theory offers two views of how this position should relate to expected returns, and they predict opposite signs. The first is a risk view. A firm that must buy allowances for a larger share of its emissions loses more when the carbon price rises or climate policy tightens, so the shortfall measures exposure to carbon pricing and transition risk. If investors demand compensation for bearing this risk, firms more exposed to carbon regulation should earn higher expected returns, as in the carbon premium of Bolton and Kacperczyk (2021), the pollution premium that Hsu et al. (2023) attribute to environmental policy risk, and the equilibrium of Pástor et al. (2021), in which green assets have low expected returns in part because they hedge climate risk.[4] The premium can also reflect the path of policy rather than realized cash flows alone: the cap tightens over time, the cost of insuring against carbon tail risk is higher for carbon-intensive firms (Ilhan et al. 2021), and carbon risk appears to be priced across markets (Bolton and Kacperczyk 2023). Under this view, high-shortfall firms should earn higher returns, the spread should load on carbon-price and systematic risk, and it should be stronger where the EU ETS is economically material to the firm.

The second is an information view. The regulatory position may carry value-relevant information about a firm's fundamentals that investors incorporate only gradually. Prices

[3] A product benchmark is the average greenhouse gas emission performance of the 10% best-performing installations in the EU that produce the product, and an installation's free allocation is the benchmark times its historical activity level, adjusted by reduction and correction factors (Martin et al. 2014). From 2013, sectors not deemed at risk of carbon leakage received a share of their benchmarked emissions for free that fell from 80% to 30% by 2020, whereas sectors at risk of carbon leakage continued to receive 100% (Fuchs 2025).

[4] Lyu et al. (2026) provide causal evidence that transition risk affects stock prices: after the U.S. Supreme Court curtailed the regulatory powers of the Environmental Protection Agency in 2022, the stock prices of firms with greater exposure to carbon risk rose relative to those of other firms.

underreact when information diffuses slowly across investors (Hong and Stein 1999), when investors' attention is limited (Hirshleifer et al. 2009; DellaVigna and Pollet 2009), and when information is costly to process (Cohen and Lou 2012). A firm's allowance position has to be assembled installation by installation from the EU registry and attributed to the listed parent, and it is not easy to interpret once assembled: free allowances leave almost no trace in financial statements, and translating a position into future earnings requires judgments about carbon prices, sector-specific allocation rules and abatement costs. Institutional investors themselves report that climate-related information is hard to quantify (Krueger et al. 2020; Ilhan et al. 2023). Under this view, low-shortfall firms should earn higher subsequent returns that risk does not explain, and expectations should adjust as the information is incorporated.

To organize these views, we present a simple two-period framework. A regulated firm's cash flow is its operating profit net of the cost of covering the share of its emissions that free allowances do not cover, and the allowance price carries a priced transition risk. All investors price risk in the same way, but part of the capital in the market does not process the firm's regulatory position and values the firm as if its position were an average one. The two views then enter the expected return as two terms with opposite signs. The risk channel rewards high-shortfall firms, in proportion to how large their carbon bill is relative to the firm. The information channel rewards low-shortfall firms, because the investors who ignore the position overstate their carbon bill and understate their profits, and the resulting underpricing is corrected when cash flows are realized. The sign of the spread therefore tells us which channel dominates, and the two channels differ in what they imply for loadings on allowance-price returns, for the materiality of the EU ETS to the firm, for carbon costs and for earnings surprises. We take these implications to the data.

The EU ETS lets us confront these views with regulator-recorded data. The European Union Transaction Log (EUTL) records, for every regulated installation and year, verified emissions and free allocation. We map installations to their listed parents through majority-

ownership chains in Orbis, assigning parents year by year, and aggregate verified emissions and allocations to the parent.[5] We exclude utilities, because electricity generators have received no free allocation since 2013, so that their shortfall is mechanically close to one. The sample contains 5,040 firm-years for 488 European firms, with 286 to 364 firms in each formation year from 2009 to 2023.

We sort firms each May into value-weighted quintiles on the shortfall of the previous compliance year and hold the portfolios for twelve months, from May 2009 to April 2024.[6] The portfolio that is long the lowest-shortfall quintile and short the highest-shortfall quintile (L−H) earns an average return of 5.4% per year with a t-statistic of 2.62, and its Fama–French five-factor alpha is 6.9% per year. The persistent component of the shortfall carries this predictability. Sorting on the persistent component yields an L−H return of 5.8% per year with a t-statistic of 2.68 and an alpha of 6.1% per year after the five factors and momentum, whereas the transitory component does not predict returns. The spread is positive in each of the three EU ETS phases that the sample covers and within terciles of vendor Scope 1 emission intensity, the persistent spread holds after excluding microcaps and within industries, and it is robust to excluding any single firm.

The evidence runs against each prediction of the risk view. The sign is the opposite of what the view predicts: firms whose emissions fall short of their allocation, the least exposed to carbon pricing, earn more. The low-shortfall portfolio is not the riskier one: its loading on allowance-price returns is insignificant, it has a lower market beta and a smaller maximum drawdown than the high-shortfall portfolio, and it does relatively well when the market falls. The spread is concentrated where the EU ETS matters least: among firms whose verified emissions are small relative to their sales, the L−H return is 4.8% per year with a t-statistic of 3.12, and among firms for which the system is material there is no

[5] Appendix C documents the mapping and its validation against an independent mapping based on legal entity identifiers and against published EUTL studies.

[6] Verified emissions for year $t-1$ must be reported by 31 March of year t, and the surrender deadline was 30 April for every compliance year in the sample (Appendix B); forming portfolios at the end of April therefore uses only information that is public at formation. Value weights use end-of-April market capitalizations, the last observation before the holding period.

spread. Finally, the net carbon costs of high-shortfall firms are small, 0.04% of sales for the highest-shortfall quintile, and the operating performance of low- and high-shortfall firms does not diverge after formation. The spread is therefore difficult to reconcile with compensation for systematic or carbon-price risk.

The evidence is consistent with the information view. In the framework, investors who value a firm at the reference position expect too little from low-shortfall firms, so these firms should beat expectations when they report earnings, and the surprise should remain when the carbon bill is negligible. Both predictions hold. In annual Fama–MacBeth regressions that control for size, book-to-market, profitability, asset growth, leverage and momentum, a lower shortfall predicts more favorable analyst forecast errors and earnings surprises, with t-statistics of 3.46 and 2.70, and the quintile means show the same ordering. Because the carbon bill is small, the surprise cannot come from a bill that analysts overlook; it reflects information in the position about the firm's business beyond the bill. The spread is also priced where investors trade on firms' cash flows. Within the largest third of firms, where the median firm is followed by 21 analysts, the value-weighted L−H return is 4.2% per year with a t-statistic of 2.36, and among covered firms the persistent spread is 5.6% per year; smaller firms and firms that no analyst follows show no spread. Within the largest tercile the equal-weighted spread is less than half the value-weighted one, so the spread rises with size even among large firms. The correction therefore occurs among firms whose cash flows investors and analysts scrutinize.[7]

Related literature and contribution. Our paper contributes to three strands of literature. First, it relates to the research on whether carbon emissions are priced in stock returns (Giglio et al. 2021), where the evidence on the sign of the premium is mixed. Some studies find that firms with higher emissions earn higher returns (Bolton and Kacperczyk 2021, 2023), others that green stocks have outperformed (Pástor et al. 2022), and others that the

[7] Appendix Tables D4 and D5 report equal-weighted portfolios, Fama–MacBeth return regressions and value-weighted results that exclude microcaps.

carbon premium depends on how emissions are measured and when they become public (Aswani et al. 2024; Zhang 2025; Atilgan et al. 2025).[8] These studies sort firms on how much carbon they emit, as reported by firms or estimated by data vendors. A related line of work studies stated commitments: Bauer et al. (2026) show that time-stamped corporate green pledges raise stock prices and predict declines in emissions. We study neither emission amounts nor stated intentions but a regulator-recorded position, audited emissions measured against a legally binding allocation, which is public at portfolio formation. Two firms with the same emissions can have opposite shortfalls. We show that this position predicts returns beyond emission intensity, and that the return pattern runs opposite to the one a carbon-risk premium implies.

Second, our paper relates to the literature on the economic consequences of carbon pricing, much of which studies the EU ETS: its effects on emissions and activity, the rents created by free allocation, and firms' compliance and trading.[9] Within this literature, the papers closest to ours study how stock prices respond to the carbon market. The response of returns to carbon-price changes varies with firms' carbon costs (Millischer et al. 2023) and with the share of verified emissions covered by free allowances (Bolton et al. 2026; Laurs et al. 2026), announcements of non-compliance lower stock prices (Laurs et al. 2026), and allocation shortfalls go with lower firm values (Clarkson et al. 2015; Dewaelheyns et al. 2023). These papers establish that the position governs a firm's exposure: when the carbon price rises, high-shortfall firms lose more. We ask a different question, whether the position is priced in expected returns, and the answer runs against the exposure logic. The firms that Bolton et al. (2026) and Laurs et al. (2026) identify as most exposed to the carbon price earn lower, not higher, subsequent returns; the predictability sits in the persistent

[8] European stocks carry a negative greenium (Alessi et al. 2021), and Chiappari et al. (2025) also use verified emissions, to build an equity index that reflects EU ETS firms' abatement efforts.

[9] The system reduced emissions without detectable losses in activity or employment (Dechezleprêtre et al. 2023; Colmer et al. 2025), free allocation combined with the pass-through of carbon costs created windfall profits (Sijm et al. 2006), firms hold permits beyond their emissions (Fuchs 2025), and firms that cut emissions sell more allowances (Akey et al. 2025). Oestreich and Tsiakas (2015) find that German firms with free allowances outperformed in the early phases.

component of the position; and it is difficult to reconcile with carbon-price risk or with a subsequent erosion of operating profits. Exposure to the carbon price and the pricing of the position are therefore two different properties of the regulatory carbon position, and the second tells us how the market processes the first.[10]

Third, our paper relates to the literature on how prices incorporate information that is slow to diffuse, receives limited attention or is costly to process and interpret (Hong and Stein 1999; Hirshleifer et al. 2009; DellaVigna and Pollet 2009; Cohen and Lou 2012; Cohen et al. 2013; Hirshleifer et al. 2013). Regulatory carbon data are a natural setting for such frictions: the information is public and audited, but scattered across installations, absent from standard financial databases and from the accounts, and its earnings implications depend on carbon prices, allocation rules and abatement technology that investors find hard to assess (Krueger et al. 2020; Ilhan et al. 2023). We find that the information predicts returns and earnings surprises, and that the spread is concentrated among large firms whose cash flows investors and analysts scrutinize. Following the standards of the anomalies literature, which shows that many predictors fail once microcaps are mitigated and returns are value-weighted (Hou et al. 2020), we also show that the persistent result holds after excluding microcaps and within industries, and that it is robust to excluding any single firm (Appendix Table D3).

The rest of the paper is organized as follows. Section 2 presents the framework and maps it to the empirical analysis. Section 3 describes the institutional setting and the data and constructs the allocation shortfall and its persistent component. Section 4 shows that the shortfall predicts the cross-section of returns, that its persistent component carries this predictability, and that it contains information beyond emission intensity. Section 5 confronts the risk and information views with the evidence on risk, materiality, carbon

[10]We also contribute a reproducible mapping from installations to listed parents with time-varying ownership. It assigns each installation-year to a single holder and listed parent, covers about 99% of emissions at the holder level, and agrees with an independent mapping based on legal entity identifiers for 99.0% of the emissions that both assign. Existing mappings use Phase 1 ownership (Jong et al. 2013), operators (Fuchs 2025) or stake-weighted attribution to listed shareholders (Laurs et al. 2026).

costs, earnings expectations and where the spread is priced. Section 6 concludes.

2 Framework

In this section, we present a simple analytical framework to illustrate the two effects that we aim to capture empirically: a risk premium for carbon-price exposure, and a correction of prices that incompletely incorporate a firm's regulatory carbon position. We then explain how we map the framework to the empirical analysis.

2.1 Analytical Framework

The agents of the model are a regulated firm and two groups of investors. The firm has a fixed regulatory position and a cash flow that depends on the carbon price. One group of investors prices the firm using its own position; the other group values the firm as if its position equalled a reference level. All investors price risk with the same pricing kernel. The market price is the capital-weighted valuation of the two groups, and the expected return is the expected cash flow relative to that price.

Firm: There are two periods, 0 and 1. The firm has verified emissions V and free allowances A, both recorded by the regulator and known at period 0. Its allocation shortfall and its period-1 cash flow are

$$S = \frac{V - A}{V}, \qquad X = \Pi - p\, S\, V, \tag{1}$$

where S is the share of emissions the firm must cover by purchasing allowances (negative for a surplus), Π is operating profit, and p is the allowance price at period 1. Because benchmarks, historical activity levels and carbon-leakage status are fixed within a trading phase, the shortfall is a persistent characteristic rather than news. The carbon bill in (1) is the mechanism in Bolton et al. (2026): a firm with a shortfall loses, and a firm with a surplus gains, when the carbon price rises. Operating profit may depend on the position

beyond the carbon bill,

$$\Pi = \bar{\Pi} - \theta\,(S - S^{\text{ref}}) + \eta, \qquad \theta \geq 0, \tag{2}$$

where S^{ref} is a reference level such as the industry average and η is a mean-zero idiosyncratic shock. We think of θ as a reduced form for the relation between a firm's benchmark treatment and activity relative to its historical base, and the quality and trajectory of the regulated business; we take no stand on its source. The constant absorbs differences in the level of profitability across firms, and θ captures only the component of future profit that covaries with the position conditional on what the market already knows. With $\theta = 0$, the position matters only through the carbon bill.

Carbon price: The allowance price carries systematic transition risk, and the pricing kernel rewards exposure to it:

$$p = \bar{p} + b\,f + \varepsilon, \qquad m = 1 + \lambda f, \qquad b \geq 0,\ \lambda \geq 0, \tag{3}$$

where f is a transition-risk factor with mean zero and unit variance, ε is a mean-zero idiosyncratic shock, and η and ε are independent of f. The riskless rate is zero. States in which the carbon price is high are states of high marginal utility, so λb is the premium per unit of carbon-price exposure.

Investors: The regulatory position is public but costly to assemble and interpret. A fraction $\mu \in [0, 1]$ of capital does not use the firm's own position and values the firm as if its shortfall were the reference level; the remaining fraction values it using the firm's own shortfall. Investors in the first group know the firm's emissions, which data vendors report, but not its allocation; the position, not the emission level, is what they do not process.

Equilibrium: The full-information value and the market price are

$$P^*(S) = \mathbb{E}[mX] = \bar{\Pi} - \theta\,(S - S^{\text{ref}}) - S\,V\,(\bar{p} + \lambda b), \tag{4}$$

$$P(S) = (1-\mu)\,P^*(S) + \mu\,P^*(S^{\text{ref}}) = P^*(S) + \mu\,(S - S^{\text{ref}})\big[\theta + V(\bar{p} + \lambda b)\big]. \tag{5}$$

At period 1 the cash flow is realized and any mispricing is corrected. The expected excess return is the expected cash flow relative to the market price, $\mathbb{E}[r(S)] = \mathbb{E}[X]/P(S) - 1$.

Proposition 1 (risk channel): With full information ($\mu = 0$), the expected excess return is

$$\mathbb{E}[r^*(S)] = \frac{\lambda\,b\,S\,V}{P^*(S)}, \tag{6}$$

which increases in the shortfall and, for a given shortfall, is proportional to the carbon weight V/P^*. Proof: See Appendix A for all proofs.

Proposition 1 says that a firm that must buy more of its emissions has cash flows that fall when the carbon price rises, and investors are compensated for the exposure; the information in the position is fully in the price and earns nothing. The proposition reproduces the comparative static of Bolton et al. (2026): a change in the carbon price moves the value of the firm in proportion to its shortfall.

Proposition 2 (information channel): With partial processing ($\mu > 0$), the expected excess return is

$$\mathbb{E}[r(S)] = \frac{\lambda\,b\,S\,V + \mu\,(S^{\text{ref}} - S)\big[\theta + V(\bar{p} + \lambda b)\big]}{P(S)}, \tag{7}$$

and, when expected returns are small relative to one, it decreases in the shortfall if and only if

$$\mu\,(\theta + \bar{p}\,V) > \lambda\,b\,(1-\mu)\,V. \tag{8}$$

Proposition 2 follows from (5): a firm whose shortfall is below the reference level is underpriced, because part of the market overstates its carbon bill and, if the position is

informative, understates its profit, and the underpricing is corrected at period 1. The two terms in (7) have opposite signs in the shortfall. Condition (8) says which dominates: the share of capital that does not process the position, weighted by what the position is worth, against the share that does, weighted by the carbon-price risk premium. The term $\mu\theta$ does not depend on the firm's emissions; if the position is informative and only partly priced, low-shortfall firms earn more even when their carbon bill is negligible.

2.2 Linking the Analytical Framework to the Empirical Analysis

We apply the ideas in this framework to the cross-section of EU ETS firms. We measure the shortfall S from the EU registry (Section 3). We do not observe the reference level S^{ref}. Our main sorts use the full cross-section without controls, which treats S^{ref} as common to all firms in a formation year; demeaning the shortfall within industries, which lets S^{ref} vary across industries, also gives positive spreads (Appendix Tables D2 and D3). We proxy for the carbon weight V/P with verified emissions relative to sales, which measures the materiality of the ETS to the firm. We proxy for the carbon-price innovation $p - \bar{p}$ with the return on EU allowance futures. In the framework, a portfolio's loading on that return measures its carbon exposure, and Proposition 1 says that a risk premium rewards this exposure rather than the position itself. We do not observe θ or μ. The framework requires only that $\mu > 0$, so that part of the market does not price the position. Whether $\theta > 0$ is an empirical question that we address through analysts' expectations. Investors who price the firm at the reference level expect earnings of $\bar{\Pi} - S^{\text{ref}} V \bar{p}$, so the expected earnings surprise of a firm with shortfall S is

$$(S^{\text{ref}} - S)\,(\theta + \bar{p}\,V), \tag{9}$$

which is positive for low-shortfall firms; when the carbon bill is negligible, a surprise remains only if $\theta > 0$. While the EU ETS provides a close empirical counterpart to this framework, the framework is more general and applies to any persistent public characteristic

that both loads on a priced factor and is costly for investors to process.

The framework organizes the empirical analysis in three ways. First, the sign of the spread identifies the dominant channel: low-shortfall firms earn higher subsequent returns if the information channel dominates and lower returns if the risk channel dominates. Because the framework treats the position as a persistent characteristic, the persistent part of the shortfall should carry the predictability, and because the investors who do not process the position still know the firm's emissions, the shortfall should predict returns beyond emission intensity. Second, the risk channel is a loading on the allowance-price factor scaled by the carbon weight. If it contributed to the spread, the spread would load on allowance returns and would be larger where the EU ETS is material to the firm. Third, the information channel implies that the earnings of low-shortfall firms surprise the market favorably, by the amount in (9). The surprise has two parts, the carbon bill and the information in the position beyond it; we gauge the first with carbon costs and subsequent profitability, and the second with analysts' forecast errors.

The framework makes no prediction about which firms, by size or analyst coverage, display the spread: the correction at period 1 is complete whatever the value of μ, whereas in the data correction within a twelve-month holding period requires investors who trade on the realized cash flows. We therefore report the size and coverage results of Section 5.4 as evidence on where the correction occurs, consistent with the information channel, rather than as implications of the framework.

3 Data and Measurement

In this section, we describe the institutional setting and the data, construct the allocation shortfall as a measure of the firm's regulatory carbon position, decompose it into a persistent and a transitory part, and present summary statistics.

3.1 Institutional Setting and Data Sources

The EU ETS caps the emissions of power plants and energy-intensive industrial installations in the European Union and the European Economic Area. Each installation reports its verified emissions for a calendar year by 31 March of the following year and must surrender an equal number of allowances (EUAs). The rules for free allocation changed across trading phases (Ellerman et al. 2016). In Phase 2 (2008–2012), national allocation plans handed out most allowances for free, largely based on historical emissions. From Phase 3 (2013–2020), auctioning became the default for electricity generation, and industrial installations received free allocation based on EU-wide product benchmarks, with more generous treatment for sectors exposed to carbon leakage (Martin et al. 2014). Phase 4 (from 2021) tightened the benchmarks and the cap. As a result, electricity generators have bought most of their allowances since 2013 (Zaklan 2023), whereas industrial installations still receive free allocation benchmarked against the most efficient installations in their sector. Appendix B gives more detail.

Installation-level data on verified emissions and free allocation come from the EUTL, in the version compiled by Abrell (2024); we use stationary installations only. The mapping from installations to listed firms proceeds in three steps, detailed in Appendix C. First, we assign each installation-year to the account holder that operated it at year end and match holders to Orbis entities through legal entity identifiers, registration numbers, Orbis identifiers and validated names. Second, we follow the chain of majority shareholders (stakes above 50%) upward to the nearest listed ancestor with a priced equity security. Because Orbis on WRDS reports current ownership links only, we correct 34 known ownership changes with dated overrides and walk past parents that were delisted before the year in question. Third, we link listed parents to Compustat Global through ISIN, SEDOL and CUSIP histories and aggregate verified emissions and allocations to the firm-year.

Monthly total returns in U.S. dollars, market capitalization and accounting data come

from Compustat Global. The sample consists of non-financial firms (excluding SIC codes 60–67 and 45) headquartered in Europe, with a European listing priced in the formation year, positive verified emissions in the prior compliance year and the standard control variables available. Factor returns are the European market, size, value, profitability, investment and momentum factors from Kenneth French's data library (Fama and French 2012, 2015); the risk-free rate is the one-month U.S. Treasury bill rate from the same source. Vendor emissions come from LSEG (formerly Refinitiv) and analysts' forecasts from I/B/E/S; we describe them where we use them.

We also exclude utilities from the sample. A firm is a utility if its Compustat SIC code is between 4900 and 4999 or its GICS sector is 55 (Utilities), using the firm's classification in the formation year and its latest available classification when the formation-year value is missing.[11] We exclude utilities for two reasons. First, utilities are regulated firms, which studies of the cross-section of stock returns commonly exclude together with financial firms. Second, and more important, under Directive 2009/29/EC electricity generators have received no free allocation since 2013: their allowances are auctioned, apart from transitional free allocation under Article 10c in some lower-income member states. The shortfall of a power generator is therefore mechanically close to one and does not vary with its position relative to a free-allocation benchmark. Appendix Table D1 reports the main results for the full sample including utilities.

3.2 The Allocation Shortfall as a Regulatory Carbon Position

For firm i and compliance year y, the allocation shortfall is

$$s_{i,y} = \frac{V_{i,y} - A_{i,y}}{V_{i,y}}, \tag{10}$$

[11]The GICS rule adds two firms that Compustat codes outside SIC 4900–4999: E.ON, a power and gas utility coded as a conglomerate (SIC 9997), and Rubis, a distributor of gas and fuel (SIC 5171).

where $V_{i,y}$ is verified emissions and $A_{i,y}$ is total free allocation, both summed over the firm's installations. Allocation includes allowances from the new entrants reserve and transitional free allocation under Article 10c of the ETS Directive. The shortfall is bounded above by one, which it reaches when a firm receives no free allocation, and is unbounded below. We winsorize it at the 1st and 99th percentiles in each formation year and do not truncate it.

The shortfall measures the firm's regulatory carbon position: the share of its emissions that it must cover by buying allowances, or, when negative, the surplus that the regulator grants it. Free allocation depends on product benchmarks, on the historical activity levels to which they are applied and on the sector's carbon-leakage status, and in Phase 2 it largely followed historical emissions (Appendix B). A low shortfall can therefore arise from producing with less carbon than the benchmark, from output below the historical activity level on which the allocation rests, or from generous regulatory treatment. We interpret the shortfall as this combination of operations and regulatory treatment. Comparisons within industries, which we report in the Appendix, hold part of the regulatory treatment fixed, but not all of it.

We decompose the shortfall into a persistent and a transitory part. The persistent part is the average of the raw shortfall over the three compliance years before formation,

$$\bar{s}_{i,y} = \frac{1}{3} \sum_{k=0}^{2} s_{i,y-k}, \tag{11}$$

computed when at least two of the three years are available.[12] The transitory part is $s_{i,y} - \bar{s}_{i,y}$. Both parts are winsorized in the same way as the level. Because benchmarks, leakage status and production technology change slowly, the persistent part captures the slow-moving regulatory position, whereas the transitory part captures year-to-year noise

[12]Requiring at least two of the three years keeps 4,879 of the 5,040 firm-years. A three-year window averages out single-year fluctuations in output and allocation while keeping the signal close to the formation date.

from output, weather and the business cycle. The persistent part has a correlation of 0.84 with the level of the shortfall, and the transitory part a correlation of 0.52. Because the two parts are correlated, we also use the residual from a cross-sectional regression of the persistent part on the transitory part within each formation year ("persistent $\perp$ transitory"), and vice versa.

3.3 Summary Statistics

Table 1 reports summary statistics. Panel A shows that the number of firms ranges from 286 to 364 per formation year. The median shortfall is negative in each formation year up to 2013 and positive from 2014 onward, and the share of firms that are short of allowances rises from 23% in 2009 to 82% in 2023. Figure 1 shows the same shift by quintile: the jump in 2014, the first formation year that uses Phase 3 allocations, reflects the move to harmonized, benchmark-based free allocation.

Panel B of Table 1 reports pooled distributions. The raw shortfall has extreme negative values, which winsorization largely removes. The median firm-year has a shortfall of 0.05, verified emissions of 70 kilotonnes and a market capitalization of USD 1.62 billion.

4 The Allocation Shortfall and the Cross-Section of Stock Returns

In this section, we show that the allocation shortfall negatively predicts value-weighted returns, that its persistent component carries this predictability, that it contains information beyond emission intensity, and that the spread is positive in every EU ETS phase.

4.1 Portfolio Sorts on the Shortfall Level

To investigate the link between the shortfall and the cross-section of stock returns, we sort firms at the end of April of each year t into quintiles on the shortfall of compliance year $t-1$, which is public by then (Section 3.1). Portfolios are value-weighted by end-of-April market capitalization and held from May of year t to April of year $t+1$. We ask whether the capital invested in regulated firms prices their regulatory position, and value weights measure the return on that capital; they also keep the many small firms in a European sample, whose prices are noisier and harder to trade, from driving the results (Hou et al. 2020). Formation years run from 2009 to 2023, so the holding period covers May 2009 to April 2024 (180 months).[13] The L−H portfolio is long the lowest-shortfall quintile and short the highest-shortfall quintile. We report alphas from the CAPM and the three-factor, five-factor and five-factor-plus-momentum models (Fama and French 1993, 2015; Carhart 1997), with Newey and West (1987) t-statistics using six lags, annualized by multiplying monthly values by 12.

Panel A of Table 2 reports the results. The L−H portfolio earns an average return of 5.4% per year with a t-statistic of 2.62. The first four quintiles earn similar average excess returns, between 10.6% and 11.4% per year, and the spread comes from the weak performance of the highest-shortfall quintile, whose five-factor alpha is negative and significant at the 5% level.

Standard factors do not explain the spread. The L−H alphas from all four factor models are between 5.7% and 6.9% per year and significant at the 5% level or better; the five-factor alpha is 6.9%, with a t-statistic of 2.65.[14] Panel B shows that decile sorts produce a similar

[13] The sample ends with the last twelve-month holding period formed under the compliance calendar in which allowances for year $t-1$ were surrendered by 30 April of year t. Directive (EU) 2023/959 moved the surrender deadline to 30 September from compliance year 2023 onward, which changes the timing of compliance relative to our end-of-April formation date; we therefore stop at the last cycle under the original calendar.

[14] Appendix Tables D2 and D3 report robustness tests: a 10% cap on portfolio weights, tercile and decile sorts, alternative ownership and allocation definitions, excluding firm-years dominated by the United Kingdom, within-industry sorts and dropping one firm at a time. Appendix Tables D4 and D5 report equal-weighted portfolios, Fama–MacBeth regressions and results that exclude microcaps. An Internet

spread, significant at the 10% level. Figure 2 plots the cumulative L−H return.

Table 3 reports the average characteristics of the quintiles. Firms in the lowest-shortfall quintile are smaller than those in the highest-shortfall quintile and have higher book-to-market ratios; their operating profitability, asset growth, leverage and momentum do not differ significantly. Emission levels do not separate the extreme quintiles either: the median firm in both emits about 35 kilotonnes, and the low-shortfall quintile is in fact the more carbon-intensive of the two, with a median of 51 tonnes of verified emissions per USD million of sales against 31 for the high-shortfall quintile (Table 3), so a low shortfall does not mean low emissions. The shortfall is therefore not a proxy for how much a firm emits.

The sector composition of the extreme quintiles differs. By the firm's dominant EU ETS activity, combustion installations account for 75% of the low-shortfall quintile's value weight and 89% of the high-shortfall quintile's. By industry, chemicals and pharmaceuticals (SIC 28) account for half of the low-shortfall quintile's value weight, against a quarter of the high-shortfall quintile's, which in turn has larger weights in communications, food and petroleum refining. Part of the L−H spread may therefore reflect industry rather than the shortfall itself; the within-industry sorts in Appendix Tables D2 and D3 address this. The low-shortfall portfolio also includes large firms for which the EU ETS is economically small and whose shortfall ratio can be extreme because verified emissions, the denominator, are small; Section 5.1 shows that the spread is in fact concentrated among such firms.

4.2 Persistent versus Transitory Shortfall

Table 4 sorts firms on the persistent and transitory parts of the shortfall. The sample requires both parts and drops 161 of the 5,040 firm-years; on this sample, the level of the shortfall yields an L−H return of 4.8% per year, somewhat below the full-sample value.

The persistent part predicts returns more strongly than the level. Panel A shows that its L−H return is 5.8% per year with a t-statistic of 2.68, and its alphas are significant at the

Appendix contains secondary tables and figures.

5% level or better in all four factor models. The transitory part does not predict returns: its L−H return is negative and insignificant.

The result holds when each part is orthogonalized to the other. The persistent part orthogonalized to the transitory part produces an L−H return of 6.7% per year with a t-statistic of 3.43, and a five-factor alpha of 7.7% per year, whereas the orthogonalized transitory part produces no spread. Panel B shows that this premium comes from the low-shortfall side. The lowest quintile of the orthogonalized sort has a five-factor alpha of 6.3% per year, significant at the 1% level, while the alphas of the other quintiles are small and insignificant.

The return predictability of the shortfall therefore comes from its slow-moving component, the firm's regulatory carbon position, and not from year-to-year fluctuations in output or allocation. Firms whose verified emissions have stayed below their free allocation over several years earn higher returns than firms that have persistently emitted more than their allocation; a surplus that arises from a single year's fluctuation does not predict returns.

4.3 Beyond Emission Intensity and Levels

If the shortfall merely restated how carbon-intensive a firm is or how much it emits, it would add little to existing measures. We examine next whether the shortfall predicts returns after controlling for vendor emissions.

We use conditional double sorts. In each formation year, we sort firms with both the shortfall and the control into terciles of the control, then form value-weighted shortfall quintiles within each tercile. We report the L−H spread within each tercile and the average of the three tercile spreads, with its five-factor alpha. For comparison, we also report the unconditional L−H spread in the same control-available sample. The controls come from LSEG: Scope 1 intensity (Scope 1 emissions over sales), Scope 1 emissions and Scope 1+2 emissions, all for the latest fiscal year ending by December of year $t-1$. LSEG emissions are available for about 150 firms per formation year, compared with about 340 in the full

sample.

Table 5 reports the results. Emission intensity does not capture the shortfall. Controlling for Scope 1 intensity, the average L−H spread across terciles is 5.0% per year with a t-statistic of 2.12, and the spread is largest among the least carbon-intensive firms. Firms with the same carbon intensity but different shortfalls therefore earn different returns. The five-factor alpha of the average spread is positive but not statistically significant. Controlling for verified EU ETS emission intensity instead of vendor intensity gives a similar average spread for the shortfall level, which is also significant at the 5% level (Appendix Table D6).

Emission levels absorb more of the spread. Controlling for Scope 1 emissions, the spread is positive in every tercile, but the average across terciles, 2.8% per year, is not statistically significant, and the same holds for Scope 1+2 emissions. The unconditional spread in the same sample is significant at the 1% level, so part of the shortfall's information overlaps with how much a firm emits.

Overall, the shortfall predicts returns after controlling for emission intensity, whether measured by vendors or from verified EU ETS emissions. It carries information about a firm's position relative to its regulatory allocation that is distinct from how carbon-intensive the firm is.

4.4 EU ETS Phases

Table 6 splits the sample by the EU ETS phase of the formation year: Phase 2 (formation years 2009–2012, 48 months), Phase 3 (2013–2020, 96 months) and Phase 4 (2021–2023, 36 months). Alphas are estimated within each subperiod.

The spread is positive in every phase for all three signals. It is largest in Phase 4, where the L−H return is about 14% per year for the shortfall level and for both persistent signals, and it is between 3% and 5% per year in Phases 2 and 3, where the shorter subperiods leave the estimates less precise.

5 Risk or Expectations?

The positive spread of Section 4 says that the information channel of Section 2 dominates the risk channel. In this section we test the two channels' remaining implications: risk and carbon-price exposure, including where the EU ETS is material; carbon costs and profitability; earnings expectations; and the firms in which the spread is priced.

5.1 Risk and Carbon-Price Exposure

If the spread were a risk premium (Proposition 1), the low-shortfall portfolio should load more on priced risks or perform worse in bad times. Panel A of Table 7 reports the factor loadings of the quintiles and of the L−H portfolio. The L−H portfolio has no loading on any of the six factors that is significant at the 5% level, and the factors explain 11% of its variance, so the factor-model alphas of Section 4.1 are close to the raw spread.

The low-shortfall portfolio is, if anything, the less risky of the two extremes. The lowest-shortfall quintile has a lower market beta, volatility and downside beta than the highest-shortfall quintile, and a much smaller maximum drawdown, 23% against 37% (Internet Appendix Table IA4). At the firm level, value-weighted betas and downside betas do not differ between the extreme quintiles, and the median low-shortfall firm is more volatile, which reflects its smaller size. If exposure to systematic risk drove the spread, the long leg should be the riskier one.

A natural candidate risk in this setting is the allowance price. Panel B of Table 7 regresses portfolio returns on the monthly return of the front EUA futures contract. The loading of the L−H portfolio on the EUA return is small, negative and insignificant, with or without the market and the other factors, and the L−H alpha after controlling for the market and the EUA return is 6.4% per year, significant at the 1% level. Conditioning on allowance-price states gives the same answer: Panel A of Table 8 shows that the L−H return does not differ significantly between months with high and low EUA returns.

The L−H portfolio also performs well when the market falls. In the 82 months with a negative market excess return, the L−H return averages 10.7% per year with a t-statistic of 2.95, and the difference from the other months is significant at the 5% level (Panel B of Table 8). This pattern is a level difference, not a convex hedge: the Henriksson and Merton (1981) regression in Panel C shows that the market slopes of the L−H portfolio in up and down markets are both small and statistically indistinguishable. Panel D reports three crisis windows, the euro crisis, the COVID-19 crash and the 2022 energy shock; the L−H return was positive in the last two and close to zero in the first. The low-shortfall portfolio does relatively well when the market does badly, which is the opposite of the payoff profile that would justify a risk premium.

The risk view also predicts that the spread should be stronger where the EU ETS is material to the firm, because carbon-price and policy shocks matter more when regulated emissions are large relative to the business. We measure materiality as verified emissions of the previous compliance year over sales. Each April, we split firms at the median of this ratio and form value-weighted shortfall terciles within each half, which keeps between 53 and 55 firms in each portfolio leg. Appendix Table D7 shows the opposite of the risk view's prediction. Among firms for which the system is less material, the L−H return on the shortfall level is 4.8% per year with a t-statistic of 3.12, and the persistent spread is significant at the 5% level. Among firms for which the system is material, neither spread differs from zero. In this baseline split, the difference between the two halves is not statistically significant.[15] This concentration is also consistent with the small net carbon costs of the firms in the extreme quintiles, to which we turn next.

Taken together, the low-shortfall portfolio is not riskier by conventional measures, it carries no exposure to the allowance price, and its outperformance is concentrated among firms for which carbon regulation matters least. These results are difficult to reconcile

[15] In no specification of Table D7 is the spread significantly larger where the system is material. The pattern holds when materiality is measured as verified emissions over the firm's total Scope 1 emissions from LSEG, and, for the persistent shortfall, when the signal is demeaned within industries. With quintiles instead of terciles within each half, the high-materiality spreads are positive but insignificant.

with compensation for systematic or carbon-price risk. Nor is the spread the realized outperformance of green stocks that Pástor et al. (2022) attribute to unexpected increases in climate concerns: the low-shortfall quintile is the more carbon-intensive of the two extremes (Table 3), the spread holds within terciles of emission intensity (Section 4.3), and it does not load on allowance-price returns.

5.2 Carbon Costs and Profitability

If the spread reflected cash flows, for instance because investors underestimated the carbon costs of high-shortfall firms, those firms should subsequently report weaker operating performance than low-shortfall firms. This test follows Atilgan et al. (2025), who find that the carbon premium partly reflects earnings surprises, and relates to evidence that firms with high carbon emission intensity have low profitability (Oestreich and Tsiakas 2024).

Panel C of Table 9 first measures the carbon costs behind the shortfall. We value each firm's net position, $V - A$, at the average EUA price of the compliance year. Almost all firms in the highest-shortfall quintile are net buyers of allowances, and almost none in the lowest-shortfall quintile, but outside the utility sector the net cost is small relative to sales: in aggregate, it equals 0.04% of the sales of the highest-shortfall quintile.

These positions do not translate into different profitability paths. Panel A reports changes in operating performance from the formation fiscal year to one and two years later. The differences between the extreme quintiles in the changes of return on assets, operating margin, return on equity and sales growth are all insignificant. Low-shortfall firms have lower operating margins at formation, but that gap neither closes nor widens afterward. Panel B confirms the result in firm-level panel regressions with SIC2-by-year fixed effects: the shortfall does not predict changes in return on assets or operating margins.

The spread therefore does not anticipate a divergence in carbon costs or profitability. High-shortfall firms pay for allowances, but the cost is small and visible at formation, and it is not followed by a relative deterioration in operating performance.

5.3 Earnings Expectations

The information channel predicts that, if expectations do not fully incorporate the regulatory position, low-shortfall firms should surprise the market favorably when they report earnings. We test this prediction with analysts' forecasts, the one set of expectations that is directly observable. Analysts are a natural place to look: the regulatory position does not appear in the income statement or the balance sheet, and its effect on earnings runs through carbon prices, allocation rules and abatement costs that are specific to each sector, so even attentive forecasters may incorporate it only partially.

We match firms to I/B/E/S and measure, for each annual earnings announcement in the holding year, the forecast error as actual earnings per share minus the last consensus forecast issued before the announcement, scaled by the stock price at the end of the preceding month, and standardized unexpected earnings (SUE) as the same difference scaled by the dispersion of forecasts (at least three estimates). Both are winsorized at the 1st and 99th percentiles. We assign each announcement to the shortfall quintile formed at the end of April of the holding year, so the shortfall is known before the consensus is compared with the outcome.

Panel B of Table 10 reports annual Fama–MacBeth regressions of the forecast error and SUE on the shortfall, which use the whole cross-section and control for size, book-to-market, profitability, asset growth, leverage and momentum. Both slopes are negative, so firms with a lower shortfall beat analysts' expectations by more. Both are significant at the 1% level, with t-statistics of 3.46 for the forecast error and 2.70 for SUE.

Panel A reports the quintile means. Forecast errors are less negative and SUE is higher in the lowest-shortfall quintile than in the highest. The quintile means are equal-weighted across announcements and do not control for size and book-to-market, which differ across quintiles (Table 3) and are themselves related to forecast errors; the regressions do. We therefore read the evidence as consistent with, rather than proof of, expectations that do not fully incorporate the information in the shortfall: low-shortfall firms subsequently beat

the earnings expectations of analysts.

5.4 Where the Spread Is Priced: Size and Analyst Coverage

The correction of prices requires investors who trade on a firm's cash flows, so we ask where such investors are present and whether the spread is priced there. Large firms are held mainly by institutional investors, who prefer large and liquid stocks (Gompers and Metrick 2001) and consider climate risks material (Krueger et al. 2020), and they are followed by more analysts. Each year, we split firms into terciles of end-of-April market capitalization or of analyst coverage, form value-weighted shortfall quintiles within each group, and compare value- and equal-weighted spreads. We measure analyst coverage as the number of analysts in the firm's latest I/B/E/S annual earnings consensus before the end of April of the formation year. In the sample, 44% of firm-years have no coverage in our I/B/E/S extract, so the low-coverage tercile consists of these firms; we also split firms into those with and without coverage.

Table 11 reports the results. The spread sits in the largest firms. For the shortfall level, the value-weighted L−H return within the largest size tercile is 4.2% per year with a t-statistic of 2.36, and its five-factor alpha is significant at the 5% level; the median firm in this tercile is followed by 21 analysts. Among mid-sized and small firms, the value-weighted L−H return is not significantly positive. For the persistent shortfall, the spread within the largest size tercile is also significant at the 5% level.

Analyst coverage gives the same picture. Among covered firms, the persistent L−H return is 5.6% per year, significant at the 5% level, and for the shortfall level the spread is significant at the 10% level and its five-factor alpha at the 5% level. Among firms that no analyst follows, neither value-weighted spread is significant.

The correction therefore occurs among firms whose cash flows investors and analysts scrutinize: the spread is priced in the largest firms and in firms that analysts follow. Within the largest tercile, the equal-weighted level spread is less than half the value-weighted

one, so the spread rises with size even among large firms. Excluding microcaps leaves the value-weighted results unchanged (Appendix Table D5); Appendix Table D4 reports equal-weighted portfolios and Fama–MacBeth regressions.

Taken together, the evidence does not support a risk or cash-flow explanation of the spread. The low-shortfall portfolio is the safer one, its outperformance is concentrated where the EU ETS is least material, and its profitability does not improve relative to the high-shortfall portfolio. Low-shortfall firms instead beat analysts' earnings expectations, and the spread is earned mainly among large, analyst-covered firms. The evidence is consistent with investors and analysts not fully incorporating the information in firms' regulatory carbon positions.

6 Conclusion

Free allocation under the EU ETS gives each regulated firm a regulatory carbon position, and the allocation shortfall measures how far the firm's verified emissions exceed or fall short of the allowances it receives for free. Among European listed firms outside the utility sector, we show that firms whose verified emissions fall persistently short of their free allocation earn higher subsequent returns than firms whose emissions persistently exceed it: the value-weighted spread on the persistent component is 5.8% per year from May 2009 to April 2024, while the transitory component has no predictive power. The spread is not explained by standard factor models, nor by carbon-price exposure or differences in carbon costs, and it is concentrated among firms for which the system is economically small. It is consistent with investors and analysts not fully incorporating the information in firms' regulatory carbon positions: controlling for firm characteristics, low-shortfall firms subsequently beat earnings expectations, and the spread is earned among the large, analyst-covered firms whose information investors process.

These results suggest that the debate over green and brown premia may be framed

too narrowly around how much firms emit. A firm's emissions relative to its regulatory allocation contain information about future returns that emission levels and intensities do not capture, and the return pattern runs opposite to the one that compensation for carbon-price risk implies. Because the allocation combines benchmarks, historical activity and carbon-leakage status, the shortfall reflects regulatory treatment as well as operations, and separating the two is a natural next step. As free allocation is phased down and the carbon border adjustment mechanism phases in, regulatory carbon positions will change, and the installation-to-parent mapping we provide makes it straightforward to study how prices incorporate them.

References

Abrell, J. (2024). Database for the European Union Transaction Log. Data documentation, version of November 6, 2024, available at euets.info.

Akey, P., I. Appel, A. Bellon, and J. Klausmann (2025). Do carbon markets undermine private climate initiatives? SSRN Working Paper 4938460, version of September 12, 2025.

Alessi, L., E. Ossola, and R. Panzica (2021). What greenium matters in the stock market? The role of greenhouse gas emissions and environmental disclosures. *Journal of Financial Stability 54*, 100869.

Anderson, B. and C. Di Maria (2011). Abatement and allocation in the pilot phase of the EU ETS. *Environmental and Resource Economics 48*(1), 83–103.

Ang, A., J. Chen, and Y. Xing (2006). Downside risk. *Review of Financial Studies 19*(4), 1191–1239.

Aswani, J., A. Raghunandan, and S. Rajgopal (2024). Are carbon emissions associated with stock returns? *Review of Finance 28*(1), 75–106.

Atilgan, Y., K. O. Demirtas, A. Edmans, and A. D. Gunaydin (2025). Does the carbon premium reflect risk or outperformance? SSRN Working Paper 4573622, revised April 2025.

Bajgar, M., G. Berlingieri, S. Calligaris, C. Criscuolo, and J. Timmis (2020). Coverage and representativeness of Orbis data. OECD Science, Technology and Industry Working Papers, OECD Publishing, Paris. doi:10.1787/c7bdaa03-en.

Bauer, M. D., D. Huber, E. Offner, and O. Wilms (2026). Corporate green pledges. *Review of Financial Studies 39*(9), 2729–2768.

Berg, F., K. Fabisik, and Z. Sautner (2020). Is history repeating itself? The (un)predictable past of ESG ratings. Finance Working Paper 708/2020, European Corporate Governance Institute.

Bolton, P. and M. Kacperczyk (2021). Do investors care about carbon risk? *Journal of Financial Economics 142*(2), 517–549.

Bolton, P. and M. Kacperczyk (2023). Global pricing of carbon-transition risk. *Journal of Finance 78*(6), 3677–3754.

Bolton, P., A. Lam, and M. Muûls (2026). Do carbon prices affect stock prices? *Journal of Financial Research 49*(2), 631–665. First published online 9 June 2025.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Chiappari, M., F. Scotti, and A. Flori (2025). Hedging financial risks with a climate index based on EU ETS firms. *Energy 320*, 135277.

Clarkson, P. M., Y. Li, M. Pinnuck, and G. D. Richardson (2015). The valuation relevance of greenhouse gas emissions under the European Union carbon emissions trading scheme. *European Accounting Review 24*(3), 551–580.

Cohen, L., K. Diether, and C. Malloy (2013). Misvaluing innovation. *Review of Financial Studies 26*(3), 635–666.

Cohen, L. and D. Lou (2012). Complicated firms. *Journal of Financial Economics 104*(2), 383–400.

Colmer, J., R. Martin, M. Muûls, and U. J. Wagner (2025). Does pricing carbon mitigate climate change? Firm-Level evidence from the European Union Emissions Trading Scheme. *Review of Economic Studies 92*(3), 1625–1660.

Dechezleprêtre, A., D. Nachtigall, and F. Venmans (2023). The joint impact of the European Union Emissions Trading System on carbon emissions and economic performance. *Journal of Environmental Economics and Management 118*, 102758.

DellaVigna, S. and J. M. Pollet (2009). Investor inattention and Friday earnings announcements. *Journal of Finance 64*(2), 709–749.

Dewaelheyns, N., F. Schoubben, K. Struyfs, and C. Van Hulle (2023). The influence of carbon risk on firm value: Evidence from the European Union Emission Trading Scheme. *Journal of Environmental Management 344*, 118293.

Ellerman, A. D. and B. K. Buchner (2007). The European Union Emissions Trading Scheme: Origins, allocation, and early results. *Review of Environmental Economics and Policy 1*(1), 66–87.

Ellerman, A. D., C. Marcantonini, and A. Zaklan (2016). The European Union Emissions Trading System: Ten years and counting. *Review of Environmental Economics and Policy 10*(1), 89–107.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fama, E. F. and K. R. French (2012). Size, value, and momentum in international stock returns. *Journal of Financial Economics 105*(3), 457–472.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics 116*(1), 1–22.

Fama, E. F. and J. D. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy 81*(3), 607–636.

Fuchs, M. (2025). Carbon permit holdings. Working paper, Copenhagen Business School, version of June 2, 2025 (SSRN Working Paper 5278572).

Giglio, S., B. Kelly, and J. Stroebel (2021). Climate finance. *Annual Review of Financial Economics 13*, 15–36.

Gompers, P. A. and A. Metrick (2001). Institutional investors and equity prices. *Quarterly Journal of Economics 116*(1), 229–259.

Henriksson, R. D. and R. C. Merton (1981). On market timing and investment performance. II. statistical procedures for evaluating forecasting skills. *Journal of Business 54*(4), 513–533.

Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental Economics and Management 59*(1), 43–56.

Hirshleifer, D., P.-H. Hsu, and D. Li (2013). Innovative efficiency and stock returns. *Journal of Financial Economics 107*(3), 632–654.

Hirshleifer, D., S. S. Lim, and S. H. Teoh (2009). Driven to distraction: Extraneous events and underreaction to earnings news. *Journal of Finance 64*(5), 2289–2325.

Hong, H. and J. C. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance 54*(6), 2143–2184.

Hou, K., C. Xue, and L. Zhang (2020). Replicating anomalies. *Review of Financial Studies 33*(5), 2019–2133.

Hsu, P.-H., K. Li, and C.-Y. Tsou (2023). The pollution premium. *Journal of Finance 78*(3), 1343–1392.

Ilhan, E., P. Krueger, Z. Sautner, and L. T. Starks (2023). Climate risk disclosure and institutional investors. *Review of Financial Studies 36*(7), 2617–2650.

Ilhan, E., Z. Sautner, and G. Vilkov (2021). Carbon tail risk. *Review of Financial Studies 34*(3), 1540–1571.

Jong, T., J. Jaraitė, A. Kažukauskas, A. Zaklan, and A. Zeitlberger (2013). Matching EU ETS accounts to historical parent companies: A technical note. Technical note, University of Groningen research database.

Krueger, P., Z. Sautner, and L. T. Starks (2020). The importance of climate risks for institutional investors. *Review of Financial Studies 33*(3), 1067–1111.

Laurs, D., P. Stork, and R. C. J. Zwinkels (2026). Pricing pollution: Asset-pricing implications of the EU Emissions Trading System. Working paper, SSRN 5145959, draft February 2026.

Lyu, X., C. Shan, and D. Y. Tang (2026). Do carbon emissions affect stock returns? Causal evidence from a Supreme Court ruling. Management Science, forthcoming (articles in advance, 2026). doi:10.1287/mnsc.2024.06732.

Martin, R., M. Muûls, L. B. de Preux, and U. J. Wagner (2014). Industry compensation under relocation risk: A firm-level analysis of the EU Emissions Trading Scheme. *American Economic Review 104*(8), 2482–2508.

Millischer, L., T. Evdokimova, and O. Fernandez (2023). The carrot and the stock: In search of stock-market incentives for decarbonization. *Energy Economics 120*, 106615.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Oestreich, A. M. and I. Tsiakas (2015). Carbon emissions and stock returns: Evidence from the EU emissions trading scheme. *Journal of Banking & Finance 58*, 294–308.

Oestreich, A. M. and I. Tsiakas (2024). Carbon emissions and firm profitability. *Journal of Sustainable Finance & Investment 14*(4), 766–786.

Pástor, Ľ., R. F. Stambaugh, and L. A. Taylor (2021). Sustainable investing in equilibrium. *Journal of Financial Economics 142*(2), 550–571.

Pástor, Ľ., R. F. Stambaugh, and L. A. Taylor (2022). Dissecting green returns. *Journal of Financial Economics 146*(2), 403–424.

Sijm, J., K. Neuhoff, and Y. Chen (2006). CO_2 cost pass-through and windfall profits in the power sector. *Climate Policy 6*(1), 49–72.

Zaklan, A. (2023). Coase and cap-and-trade: Evidence on the independence property from the European carbon market. *American Economic Journal: Economic Policy 15*(2), 526–558.

Zhang, S. (2025). Carbon returns across the globe. *Journal of Finance 80*(1), 615–645.

Table 1: Summary Statistics

This table reports summary statistics for the sort sample: non-financial European firms outside the utility sector (SIC 4900–4999 or GICS sector 55) with positive verified EU ETS emissions in the prior compliance year and the required control variables, formation years 2009–2023 (5,040 firm-years, 488 firms). Formation year t uses the shortfall of compliance year $t-1$. The shortfall is $(V-A)/V$, where V is verified emissions and A total free allocation. Panel A reports, by formation year, the number of firms, their total end-of-April market capitalization, verified emissions and free allocation, the median shortfall and the share of firms with a positive shortfall. Panel B pools all firm-years; the winsorized shortfall is winsorized at the 1st and 99th percentiles within each formation year. Data: EUTL, Orbis and Compustat Global.

Formation year	Firms	Market cap (USD bn)	Verified (Mt)	Allocated (Mt)	Median shortfall	Share short
Panel A: Firms by formation year						
2009	338	2,452	456.0	493.5	−0.12	0.23
2010	348	3,305	386.2	512.0	−0.29	0.14
2011	350	4,199	417.7	518.1	−0.23	0.19
2012	357	3,703	408.0	518.7	−0.25	0.18
2013	354	4,236	392.8	525.7	−0.34	0.14
2014	364	5,024	429.8	406.1	0.14	0.62
2015	361	4,768	425.2	390.7	0.14	0.65
2016	343	4,323	426.9	375.8	0.18	0.66
2017	339	4,730	440.8	371.9	0.21	0.70
2018	334	5,250	450.5	368.6	0.25	0.75
2019	325	4,842	443.6	355.0	0.25	0.77
2020	324	4,117	439.5	346.4	0.26	0.75
2021	325	5,533	415.1	341.4	0.23	0.72
2022	292	5,121	412.6	296.9	0.35	0.85
2023	286	5,733	392.4	292.9	0.36	0.82

	N	Mean	Std	P_{10}	P_{25}	Median	P_{75}	P_{90}
Panel B: Distribution over all firm-years								
Shortfall (raw)	5,040	−1.49	38.19	−0.86	−0.27	0.05	0.33	0.63
Shortfall (winsorized)	5,040	−0.16	1.17	−0.86	−0.27	0.05	0.33	0.63
Verified (kt)	5,040	1,257	4,357	5	15	70	485	2,676
Allocated (kt)	5,040	1,213	4,954	4	13	66	478	2,420
Market cap (USD m)	5,040	13,360	33,596	49	273	1,624	8,744	35,243
B/M	5,040	1.05	6.79	0.26	0.44	0.78	1.32	2.25

Table 2: Portfolio Sorts: Excess Returns and Alphas (% per Year)

This table reports average excess returns and alphas (% per year) for value-weighted portfolios sorted on the allocation shortfall. Returns and alphas are annualized by multiplying monthly values by 12. At the end of April of each year from 2009 to 2023, firms are sorted into quintiles (Panel A) or deciles (Panel B) on the shortfall of the previous compliance year and held from May to April. Portfolios are value-weighted by end-of-April market capitalization. Ret is the average excess return over the one-month Treasury bill rate; α is the intercept from time-series regressions on the CAPM, the Fama–French three- and five-factor models and the five-factor model plus momentum (European factors). L−H is the lowest-shortfall portfolio minus the highest-shortfall portfolio. N is the average number of firms. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024 (180 months).

	Mean		CAPM		FF3		FF5		FF5+MOM		
	Ret	t	α	t	α	t	α	t	α	t	N
Panel A: Quintile sort, value-weighted											
Q1 (Low)	11.44	2.87	4.44	2.33	4.15	2.13	3.73	1.66	2.69	1.11	66
Q2	10.92	2.29	1.95	1.46	2.16	1.72	0.89	0.65	1.42	0.86	67
Q3	11.29	2.76	3.61	1.92	4.29	2.61	2.42	1.36	2.41	1.18	67
Q4	10.63	2.86	3.09	1.80	3.08	1.75	2.06	1.13	2.25	1.22	67
Q5 (High)	6.04	1.46	−1.75	−1.13	−1.53	−1.00	−3.12	−2.48	−3.46	−2.41	67
L−H	5.40	2.62	6.19	2.77	5.68	2.54	6.86	2.65	6.15	2.17	
Panel B: Decile sort, value-weighted											
D1 (Low)	13.72	3.42	6.87	2.90	6.41	2.48	5.86	2.48	3.90	1.62	33
D2	10.57	2.49	3.16	1.47	2.99	1.42	3.04	1.06	2.81	0.89	33
D3	12.31	2.41	3.13	1.48	3.27	1.58	3.19	1.45	2.79	1.19	33
D4	11.30	2.35	2.43	1.37	2.72	1.58	0.77	0.39	1.87	0.86	33
D5	12.31	3.24	4.95	2.62	5.63	3.45	3.78	1.97	2.42	1.09	34
D6	9.82	2.00	1.45	0.58	2.20	0.98	0.06	0.03	1.87	0.82	33
D7	11.43	2.69	3.21	1.35	3.30	1.36	2.34	0.92	2.85	1.14	33
D8	9.65	2.71	2.63	1.48	2.59	1.50	1.57	0.75	1.40	0.60	34
D9	3.67	0.79	−4.93	−2.80	−4.89	−2.70	−6.32	−3.58	−6.76	−3.34	33
D10 (High)	8.14	2.09	1.06	0.51	1.65	0.82	0.63	0.34	0.49	0.25	34
L−H	5.58	1.86	5.80	1.83	4.76	1.40	5.24	1.72	3.41	1.06	

Table 3: Quintile Portfolio Characteristics

Panel A reports time-series averages of the cross-sectional medians (unless labeled mean or value-weighted) of firm characteristics at formation for shortfall-sorted quintiles, formation years 2009–2023. Q1−Q5 is the difference between the extreme quintiles, and t its t-statistic across the 15 formation years. Operating profitability is operating income after depreciation over assets, asset growth is the growth in total assets, Capex/AT is capital expenditure over assets, leverage is total debt over assets, $\text{Mom}_{12\text{-}2}$ is the return from month −12 to −2, and carbon intensity is verified emissions in tonnes per USD million of sales (aggregate: total emissions over total sales of the quintile). Panel B reports the average value weights of Q1 and Q5 by the firm's dominant EU ETS activity and by SIC2 industry (the ten industries with the largest weight in either quintile).

	Q1 (Low)	Q2	Q3	Q4	Q5 (High)	Q1−Q5	t
Panel A: Firm characteristics							
Shortfall (mean)	−1.46	−0.15	0.07	0.25	0.54	−2.01	−14.63
Shortfall (median)	−0.77	−0.15	0.08	0.25	0.48	−1.25	−16.30
Shortfall (value-weighted)	−1.21	−0.14	0.07	0.25	0.59	−1.80	−13.96
Market cap (USD m)	1,129	1,540	1,809	2,419	2,641	−1,512	−2.34
B/M	0.91	0.88	0.80	0.72	0.73	0.18	4.75
Operating profitability	0.056	0.059	0.065	0.069	0.060	−0.004	−0.85
Asset growth	0.027	0.025	0.040	0.045	0.040	−0.013	−1.60
Capex/AT	0.039	0.035	0.044	0.047	0.040	−0.001	−0.24
Leverage	0.239	0.253	0.236	0.233	0.247	−0.009	−0.53
$\text{Mom}_{12\text{-}2}$	0.037	0.012	0.024	0.031	0.017	0.019	1.50
Carbon intensity (median)	51.3	91.1	88.7	78.8	30.7	20.6	3.00
Carbon intensity (aggregate)	132.8	152.0	84.0	87.8	53.6	79.3	3.19
Verified (kt)	34.8	125.9	159.8	122.8	35.3	−0.5	−0.08
N (firms per year)	66.7	67.3	67.2	67.3	67.5	−0.8	−7.48

EU ETS activity	Q1	Q5	SIC2 industry	Q1	Q5
Panel B: Value weights by sector					
Combustion	0.753	0.895	28 Chemicals, pharma	0.504	0.244
Pulp and paper	0.075	0.003	48 Communications	0.015	0.154
Iron and steel	0.055	0.017	37 Transport equip.	0.127	0.119
Cement and lime	0.038	0.001	20 Food	0.026	0.124
Non-ferrous metals	0.035	0.005	29 Petroleum refining	0.000	0.104
Chemicals	0.034	0.016	33 Primary metals	0.067	0.013
Glass, ceramics	0.010	0.004	26 Paper	0.062	0.002
Refining	0.001	0.042	21 Tobacco	0.018	0.060
Other	0.000	0.018	13 Oil and gas extr.	0.001	0.051

Table 4: Persistence Decomposition: Persistent versus Transitory Shortfall

This table reports value-weighted quintile sorts on the components of the shortfall. The persistent shortfall is the mean of the raw shortfall over the three compliance years before formation (at least two available); the transitory shortfall is the level minus the persistent shortfall; both are winsorized at 1/99 within formation year. "$\perp$" denotes the residual from a cross-sectional regression on the other component within formation year. The sample requires both components (4,879 firm-years). Panel A reports L−H spreads and alphas (% per year) and the average excess returns of Q1 and Q5; Panel B reports the quintiles of the persistent $\perp$ transitory sort. Returns and alphas are annualized by multiplying monthly values by 12. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024.

	L−H		CAPM		FF3		FF5		FF5+MOM		Q1	Q5
Signal	Ret	t	α	t	α	t	α	t	α	t	Ret	Ret
Panel A: L−H spreads by signal												
Level	4.75	2.40	5.32	2.62	4.88	2.27	5.09	2.36	4.51	1.92	11.48	6.73
Persistent	5.76	2.68	6.78	2.91	6.38	2.56	7.57	3.21	6.09	2.48	13.57	7.82
Transitory	−1.44	−0.59	−0.96	−0.37	−0.70	−0.27	0.43	0.14	0.38	0.11	6.88	8.31
Persistent $\perp$ transitory	6.74	3.43	7.70	3.50	7.21	3.13	7.75	3.38	6.91	2.94	13.66	6.92
Transitory $\perp$ persistent	−1.81	−0.79	−1.29	−0.51	−1.11	−0.44	0.37	0.13	−0.56	−0.16	7.13	8.94

	Mean		CAPM		FF3		FF5		FF5+MOM		
	Ret	t	α	t	α	t	α	t	α	t	N
Panel B: Quintiles sorted on the persistent shortfall $\perp$ transitory											
Q1 (Low)	13.66	3.53	7.02	3.78	6.82	3.50	6.32	3.05	5.55	2.49	64
Q2	11.10	2.55	2.50	1.76	2.55	1.74	1.32	0.85	1.79	0.92	64
Q3	9.22	2.12	1.27	0.68	1.70	1.03	−0.09	−0.05	0.44	0.27	64
Q4	9.73	2.40	1.87	1.21	2.11	1.45	1.05	0.57	0.24	0.13	64
Q5 (High)	6.92	1.84	−0.68	−0.43	−0.39	−0.24	−1.43	−1.01	−1.36	−0.84	65
L−H	6.74	3.43	7.70	3.50	7.21	3.13	7.75	3.38	6.91	2.94	

Table 5: Double Sorts on the Allocation Shortfall, Controlling for Emission Intensity and Levels

This table reports conditional double sorts. At the end of April of each year from 2009 to 2023, firms with both the shortfall and the control are sorted into terciles of the control (T1 = lowest), and then into value-weighted quintiles on the shortfall within each tercile (at least three firms per quintile). L−H is the lowest- minus the highest-shortfall quintile within the tercile (% per year). Returns and alphas are annualized by multiplying monthly values by 12. Average is the monthly mean of the three tercile L−H returns; FF5 α is its Fama–French five-factor alpha. Unconditional is the one-way value-weighted L−H spread in the same control-available sample. Scope 1 and Scope 2 emissions are from LSEG for the latest fiscal year ending by December of year $t-1$; Scope 1 intensity is Scope 1 emissions over sales; Scope 1+2 requires both. The average number of firms per formation year is about 150 (146 for Scope 1+2). Double sorts on the ESG score, the environmental-pillar score and verified emission intensity are in Appendix Table D6; the size tercile sort is in Table 11. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024 (180 months).

	T1 (Low)		T2		T3 (High)		Average		Average		Unconditional	
Control	L−H	t	L−H	t	L−H	t	L−H	t	FF5 α	t	L−H	t
Scope 1 intensity	8.18	2.91	3.44	0.83	3.49	0.88	5.04	2.12	3.93	1.64	6.32	3.04
Scope 1 emissions	3.66	1.07	1.07	0.40	3.65	0.81	2.79	1.49	0.87	0.45	6.32	3.04
Scope 1+2 emissions	3.93	1.09	3.23	1.10	2.11	0.41	3.09	1.41	2.06	0.88	5.69	2.72

Table 6: Subperiod Analysis by EU ETS Phase

This table reports value-weighted quintile L−H spreads (% per year) by the EU ETS phase of the formation year. Returns and alphas are annualized by multiplying monthly values by 12. Phase 2 covers formation years 2009–2012 (holding months May 2009–April 2013), Phase 3 covers 2013–2020 (May 2013–April 2021) and Phase 4 covers 2021–2023 (May 2021–April 2024). Alphas are estimated within each subperiod. t-statistics use Newey-West standard errors with 6 lags.

	L−H		CAPM		FF5		FF5+MOM		
	Ret	t	α	t	α	t	α	t	Months
Panel A: Shortfall level									
Full sample (2009–2023)	5.40	2.62	6.19	2.77	6.86	2.65	6.15	2.17	180
Phase 2 (2009–2012)	3.45	1.45	3.92	1.51	1.95	0.63	1.44	0.49	48
Phase 3 (2013–2020)	3.18	0.96	4.98	1.41	4.16	1.24	4.29	1.18	96
Phase 4 (2021–2023)	13.90	3.54	13.85	3.57	15.09	4.32	14.14	4.09	36
Panel B: Persistent shortfall									
Full sample (2009–2023)	5.76	2.68	6.78	2.91	7.57	3.21	6.09	2.48	180
Phase 2 (2009–2012)	4.98	1.63	5.05	1.45	4.32	1.22	3.61	1.14	48
Phase 3 (2013–2020)	3.20	1.07	4.98	1.53	4.78	1.55	4.01	1.33	96
Phase 4 (2021–2023)	13.61	2.40	13.74	2.38	16.07	3.62	14.19	2.78	36
Panel C: Persistent shortfall $\perp$ transitory									
Full sample (2009–2023)	6.74	3.43	7.70	3.50	7.75	3.38	6.91	2.94	180
Phase 2 (2009–2012)	4.31	1.63	5.14	1.78	4.60	1.68	4.34	1.74	48
Phase 3 (2013–2020)	5.29	1.95	6.54	2.11	5.95	2.00	5.71	1.96	96
Phase 4 (2021–2023)	13.87	2.55	13.99	2.46	18.20	4.34	15.40	2.99	36

Table 7: Factor Loadings and Allowance-Price Exposure

Panel A reports time-series regressions of quintile excess returns and the L−H return on the European Fama–French five factors and momentum; α is in % per year. Returns and alphas are annualized by multiplying monthly values by 12. Panel B regresses the same returns on the monthly EUA futures return alone, with the market factor, and with the five factors and momentum. The EUA return compounds within-contract daily settlement returns of the front ICE EUA futures continuation, with roll days dropped. The L−H column of Panel B is the difference between the Q1 and Q5 EUA loadings. t-statistics (in parentheses) use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024 (180 months).

	Q1 (Low)	Q2	Q3	Q4	Q5 (High)	L−H
Panel A: Five-factor-plus-momentum loadings						
α	2.69	1.42	2.41	2.25	−3.46	6.15
	(1.11)	(0.86)	(1.18)	(1.22)	(−2.41)	(2.17)
MKT	0.92	1.04	0.89	0.94	0.96	−0.04
	(23.07)	(29.29)	(23.81)	(31.53)	(30.72)	(−0.71)
SMB	−0.12	−0.13	−0.14	−0.20	−0.23	0.11
	(−0.92)	(−1.74)	(−1.16)	(−2.45)	(−3.45)	(0.68)
HML	−0.30	0.10	0.34	−0.17	0.04	−0.33
	(−1.86)	(0.85)	(1.75)	(−1.92)	(0.39)	(−1.80)
RMW	0.06	0.32	0.45	0.23	0.37	−0.32
	(0.26)	(2.15)	(2.11)	(1.87)	(3.35)	(−1.13)
CMA	0.41	0.16	0.34	0.42	0.40	0.02
	(2.99)	(1.44)	(1.99)	(3.14)	(2.47)	(0.10)
WML	0.11	−0.05	0.00	−0.02	0.03	0.07
	(1.87)	(−0.82)	(0.01)	(−0.46)	(0.64)	(1.09)
R^2	0.82	0.92	0.84	0.90	0.90	0.11
Panel B: Loadings on EUA futures returns						
Regressors: EUA only						
b_{EUA}	0.044	0.082	0.065	0.059	0.066	−0.022
	(1.54)	(1.89)	(1.70)	(1.66)	(1.52)	(−1.02)
α	10.70	9.54	10.21	9.64	4.94	5.76
	(2.59)	(2.00)	(2.48)	(2.60)	(1.17)	(2.84)
Regressors: EUA and MKT						
b_{EUA}	−0.017	0.006	−0.001	−0.006	−0.001	−0.016
	(−1.49)	(0.54)	(−0.06)	(−0.65)	(−0.07)	(−0.85)
α	4.64	1.88	3.62	3.16	−1.74	6.38
	(2.51)	(1.38)	(1.94)	(1.83)	(−1.10)	(2.88)
Regressors: EUA and FF5+MOM						
b_{EUA}	−0.002	0.007	−0.001	0.004	0.009	−0.010
	(−0.19)	(0.65)	(−0.12)	(0.64)	(0.73)	(−0.65)
α	2.72	1.31	2.44	2.18	−3.60	6.32
	(1.12)	(0.80)	(1.18)	(1.19)	(−2.48)	(2.25)

Table 8: Conditional Performance and Down Markets

Panel A reports mean returns (% per year) by tercile of the monthly EUA futures return and by the sign of the European market excess return (MKT), from regressions on state dummies; difference rows regress on a dummy for the second state. Panel B compares the L−H return in months with MKT < 0 or in the bottom MKT tercile with the remaining months, for the level and persistent signals. Panel C estimates $\mathrm{L} - \mathrm{H}_t = \alpha + b_{\mathrm{up}} \max(\mathrm{MKT}_t, 0) + b_{\mathrm{down}} \min(\mathrm{MKT}_t, 0) + \varepsilon_t$ (Henriksson and Merton 1981); $b_{\mathrm{up}} - b_{\mathrm{down}} < 0$ would indicate that L−H gains relative to the market in down markets beyond a linear exposure. t-statistics in Panels A–C use Newey-West standard errors with 6 lags. Panel D reports the mean L−H return with its iid t-statistic and the compounded L−H and market excess returns (%) over three crisis windows. Mean returns and alphas are annualized by multiplying monthly values by 12; compounded returns are not annualized. Sample: May 2009–April 2024.

	L−H		Q1 excess		Q5 excess		
	Ret	t	Ret	t	Ret	t	Months
Panel A: Returns by EUA-return tercile and by market state							
Low EUA tercile	6.38	1.63	3.17	0.45	−3.21	−0.45	60
Middle EUA tercile	9.66	2.56	10.09	1.42	0.43	0.06	60
High EUA tercile	0.15	0.04	21.06	3.15	20.91	2.54	60
High − low	−6.23	−1.00	17.89	1.91	24.11	1.94	120
Market down (MKT < 0)	10.65	2.95	−31.89	−8.74	−42.54	−9.56	82
Market up (MKT ≥ 0)	1.00	0.38	47.69	11.73	46.69	10.93	98
Up − down	−9.65	−2.09	79.58	13.84	89.23	14.08	180

		Down months		Other months		Difference	
Signal	Down month	Ret	t	Ret	t	Ret	t
Panel B: L−H in down months							
Level	MKT < 0	10.65	2.95	1.00	0.38	9.65	2.09
Level	Bottom tercile	13.48	2.75	1.25	0.48	12.23	2.01
Persistent	MKT < 0	12.64	3.27	0.00	0.00	12.63	2.55
Persistent	Bottom tercile	14.68	3.06	1.19	0.48	13.49	2.41

Signal	α	t	b_{up}	t	b_{down}	t	$b_{\mathrm{up}} - b_{\mathrm{down}}$	t	R^2
Panel C: Henriksson-Merton regression of L−H									
Level	5.46	1.15	−0.079	−0.88	−0.109	−0.70	0.030	0.14	0.034
Persistent	9.78	2.20	−0.176	−2.12	−0.053	−0.42	−0.123	−0.69	0.061

Signal	Window	Months	L−H	t	Cum. L−H	Cum. MKT	L−H> 0
Panel D: Crisis windows							
Level	Euro crisis (2011:07–2012:06)	12	−0.83	−0.11	−1.07	−18.05	6
Level	COVID-19 (2020:02–2020:04)	3	64.41	1.40	16.52	−17.88	2
Level	Energy shock (2022:02–2022:10)	9	10.50	1.12	7.93	−22.54	6
Persistent	Euro crisis (2011:07–2012:06)	12	−2.40	−0.33	−2.60	−18.05	5
Persistent	COVID-19 (2020:02–2020:04)	3	36.86	0.82	9.07	−17.88	2
Persistent	Energy shock (2022:02–2022:10)	9	13.03	1.47	10.00	−22.54	7

Table 9: Carbon Costs and Future Profitability

Panel A reports equal-weighted means within quintile and formation year, averaged over formation years, of changes from the formation fiscal year t to $t+1$ and $t+2$ (winsorized at 1/99 within formation year), in percentage points; the last two columns report levels at t (%). ROA is operating income after depreciation over assets, margin is operating income over sales, ROE is income before extraordinary items over lagged book equity, and sales growth is in local currency. t(Q1−Q5) is computed across the 15 formation years (Newey-West, lags equal to the horizon). Panel B reports panel regressions of future changes on the shortfall and the carbon burden (net carbon cost over sales), both standardized within formation year, with controls for log market capitalization, B/M, leverage and ROA_t and SIC2 × formation-year fixed effects; coefficients are in percentage points per standard deviation, with t-statistics clustered by firm in parentheses. Panel C values the net position $V - A$ at the average EUA price of the compliance year, converted to U.S. dollars (negative values are surpluses), in % of sales or of positive operating income; Aggregate is the quintile's total cost over its total sales (or operating income); Net cost counts only shortfalls; Net buyers is the share of firms with $V > A$.

	ΔROA		ΔMargin		ΔROE		Sales growth		Level at t	
	$t+1$	$t+2$	$t+1$	$t+2$	$t+1$	$t+2$	$t+1$	$t+2$	ROA	Margin
Panel A: Changes in operating performance by quintile										
Q1 (Low)	−0.02	0.11	0.14	0.16	−0.88	−0.53	4.45	10.77	5.57	5.69
Q2	−0.13	−0.16	−0.28	−0.12	−1.09	−1.83	4.71	10.37	5.97	7.30
Q3	0.21	0.23	0.12	0.20	−0.83	0.07	5.65	12.56	6.43	7.44
Q4	−0.20	−0.08	−0.16	−0.01	−1.04	−0.06	5.84	12.65	7.26	9.74
Q5 (High)	−0.24	−0.24	−0.28	−0.40	−1.16	−1.20	5.99	14.29	6.43	8.25
Q1−Q5	0.22	0.35	0.42	0.56	0.28	0.67	−1.53	−3.52	−0.86	−2.56
t(Q1−Q5)	0.63	0.60	1.09	0.98	0.37	0.53	−1.21	−1.52	−1.26	−4.49

	ΔROA_{t+1}	ΔROA_{t+2}	$\Delta\text{Margin}_{t+1}$	$\Delta\text{Margin}_{t+2}$
Panel B: Panel regressions of future changes in profitability				
Specification (1)				
Shortfall (z)	0.073	0.008	0.086	−0.102
	(1.02)	(0.09)	(0.58)	(−0.61)
Specification (2)				
Carbon burden (z)	0.041	−0.031	−0.132	0.051
	(0.42)	(−0.27)	(−0.67)	(0.22)
Specification (3)				
Shortfall (z)	0.068	0.017	0.141	−0.130
	(0.91)	(0.21)	(0.94)	(−0.80)
Carbon burden (z)	0.021	−0.036	−0.173	0.090
	(0.20)	(−0.30)	(−0.85)	(0.37)
Observations in (1)	4,648	4,868	4,638	4,860
Firms in (1)	460	477	460	477
R^2 in (1)	0.40	0.47	0.37	0.33

	Burden / sales			Burden / operating income		Net
	Median	Aggregate	Net cost	Median	Aggregate	buyers
Panel C: Carbon cost burden at formation						
Q1 (Low)	−0.08	−0.07	0.00	−0.78	−0.65	0.02
Q2	0.00	0.00	0.04	0.09	0.04	0.41
Q3	0.04	0.03	0.04	0.50	0.24	0.67
Q4	0.04	0.07	0.08	0.35	0.47	0.68
Q5 (High)	0.01	0.04	0.04	0.15	0.32	0.95

Table 10: Analyst Forecast Errors and Earnings Surprises

Panel A reports the average analyst forecast error and standardized unexpected earnings (SUE) of annual earnings announcements by shortfall quintile. The forecast error is actual earnings per share minus the last I/B/E/S consensus (FY1 mean) issued within 100 days before the announcement, scaled by the stock price at the end of the preceding month (in % of price); SUE is the same difference scaled by the standard deviation of forecasts (at least three estimates). Both are winsorized at 1/99. Each announcement is assigned to the value-weighted-sort quintile formed at the end of April of the holding year in which it falls. Means are equal-weighted within quintile and formation year and averaged over the 15 formation years (2009–2023); t-statistics are computed across formation years; N is the number of announcements. Panel B reports slopes from annual Fama–MacBeth regressions of the forecast error and SUE on the shortfall (truncated at −1) with controls for log market capitalization, B/M, operating profitability, asset growth, leverage and momentum; t-statistics use Newey-West standard errors with 2 lags, N is the average number of firms per year and Years the number of annual cross-sections. A negative slope corresponds to more favorable surprises for low-shortfall firms.

	Forecast error (% of price)			SUE		
	Mean	t	N	Mean	t	N
Panel A: Forecast errors and SUE by shortfall quintile						
Q1 (Low)	−0.34	−1.48	499	0.25	3.07	417
Q2	−0.26	−1.04	554	0.18	1.79	482
Q3	−0.62	−2.50	590	0.13	1.41	523
Q4	−0.27	−1.46	541	0.23	2.78	483
Q5 (High)	−0.50	−2.04	532	0.12	1.12	453
Q1−Q5	0.17	0.77	2,716	0.14	1.18	2,358

Dependent variable	Slope	t	N	Years
Panel B: Fama–MacBeth regressions on the shortfall				
Forecast error (%)	−0.478	−3.46	180	15
SUE	−0.210	−2.70	156	15

Table 11: The Shortfall Premium by Size and Analyst Coverage

At the end of April of each year, firms are split into terciles of end-of-April market capitalization, into terciles of analyst coverage, or into firms with and without analyst coverage, and then sorted into shortfall quintiles within each group (at least three firms per quintile). Coverage is the number of analysts in the firm's latest I/B/E/S annual (FY1) earnings consensus issued in the twelve months up to the end of April of the formation year; firms without a consensus in the I/B/E/S extract have zero coverage (44% of firm-years), and coverage terciles use the within-year percentile rank with ties averaged, so all zero-coverage firms form the low group. Firms is the average number of firms per formation year; Analysts is the median coverage. L−H is the lowest- minus the highest-shortfall quintile (% per year), value-weighted (VW) or equal-weighted (EW) within the group; FF5 α is its Fama–French five-factor alpha. Returns and alphas are annualized by multiplying monthly values by 12. Panel B uses the persistent shortfall (three-year mean). t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024 (180 months).

			Value-weighted				Equal-weighted			
Group	Firms	Analysts	L−H	t	FF5 α	t	L−H	t	FF5 α	t
Panel A: Shortfall level										
Small size	112	0	−6.21	−1.87	−8.84	−2.51	−0.88	−0.20	0.42	0.09
Medium size	112	4	0.91	0.29	−0.33	−0.11	1.46	0.48	0.50	0.18
Large size	112	21	4.15	2.36	4.85	2.46	1.75	0.80	1.03	0.41
Low coverage	149	0	1.77	0.43	2.02	0.43	1.25	0.47	2.50	0.81
Medium coverage	75	5	2.94	0.81	1.95	0.60	1.31	0.39	1.77	0.53
High coverage	112	22	2.74	1.08	2.70	0.96	0.18	0.07	−1.85	−0.71
No analysts	149	0	1.77	0.43	2.02	0.43	1.25	0.47	2.50	0.81
Covered	187	15	4.45	1.86	5.20	1.99	0.22	0.10	0.02	0.01
Panel B: Persistent shortfall										
Small size	108	0	−3.90	−1.19	−8.09	−2.69	−1.07	−0.29	−2.06	−0.46
Medium size	108	5	−0.21	−0.08	−2.48	−0.84	1.02	0.38	−0.08	−0.03
Large size	109	22	4.24	2.40	4.30	2.18	3.77	1.89	3.25	1.46
Low coverage	142	0	2.37	0.68	4.57	1.07	4.83	1.93	5.76	1.97
Medium coverage	74	5	2.26	0.62	2.50	0.70	0.08	0.03	0.47	0.14
High coverage	109	22	4.12	1.66	4.57	1.67	0.92	0.36	−0.90	−0.34
No analysts	142	0	2.37	0.68	4.57	1.07	4.83	1.93	5.76	1.97
Covered	183	15	5.63	2.18	5.64	1.98	1.27	0.63	0.33	0.15

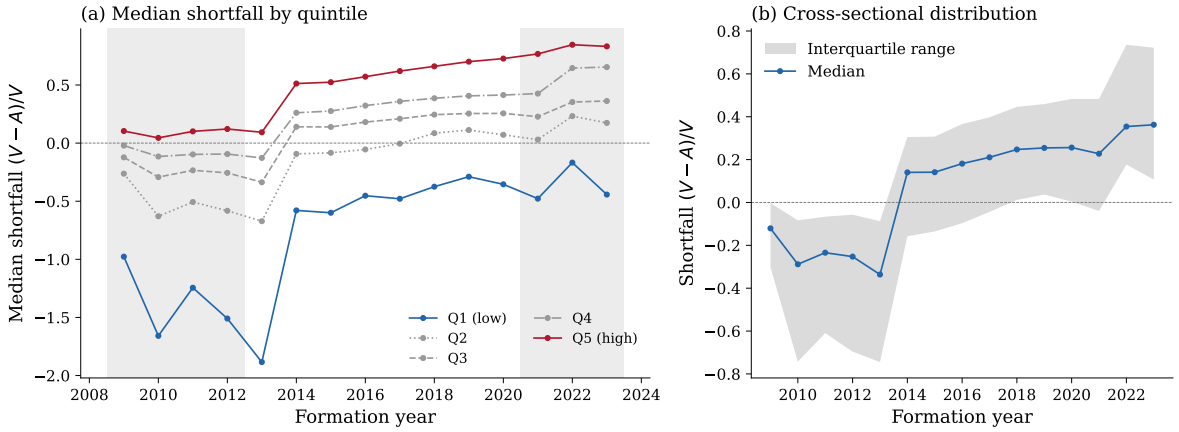


Figure 1: **Allocation Shortfall over Time.** Panel (a) plots the median allocation shortfall $(V - A)/V$ of each value-weighted quintile by formation year; Panel (b) plots the cross-sectional median and interquartile range of the winsorized shortfall. Shaded areas mark formation years in Phase 2 (2009–2012) and Phase 4 (2021–2023). The quintiles are those of Table 2.

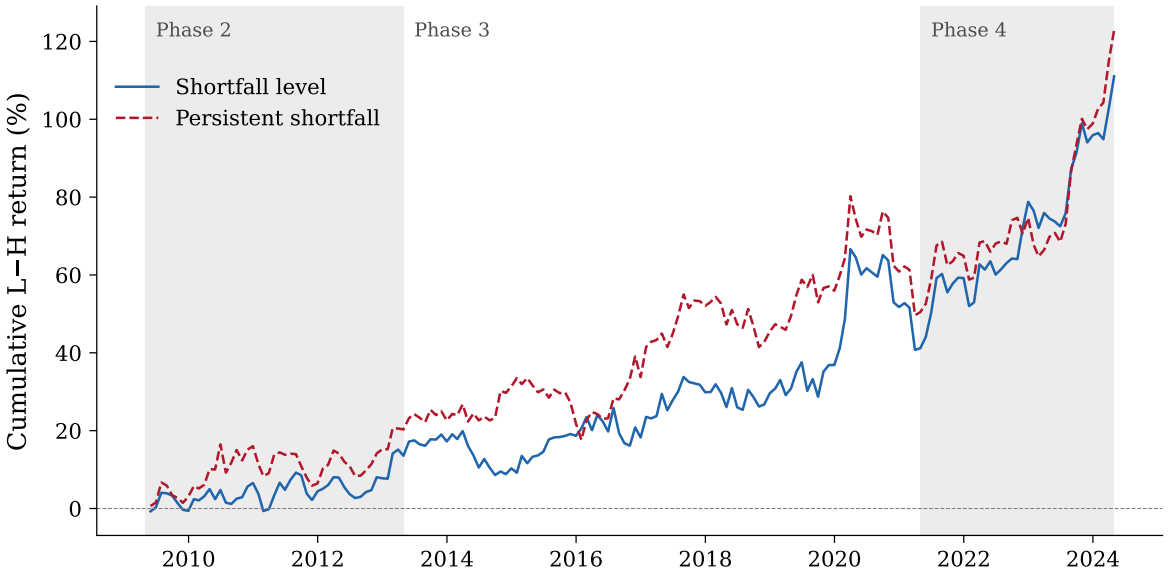


Figure 2: **Cumulative L−H Returns (VW Quintile).** This figure plots the cumulative (compounded) return of the value-weighted L−H quintile portfolio sorted on the shortfall level (solid) and on the persistent shortfall (dashed), from May 2009 to April 2024. Background shading indicates the holding periods of EU ETS Phase 2 (formation years 2009–2012) and Phase 4 (2021–2023).

Appendix A Proofs

Full-information value. Using $\mathbb{E}[f] = 0$, $\mathbb{E}[f^2] = 1$, $\mathbb{E}[\eta] = \mathbb{E}[\varepsilon] = 0$ and the independence of ε and η from f,

$$\begin{aligned}\mathbb{E}[mX] &= \mathbb{E}\big[(1+\lambda f)(\Pi - pSV)\big] = \bar{\Pi} - \theta(S - S^{\text{ref}}) - SV\,\mathbb{E}\big[(1+\lambda f)(\bar{p} + bf + \varepsilon)\big] \\ &= \bar{\Pi} - \theta(S - S^{\text{ref}}) - SV(\bar{p} + \lambda b),\end{aligned}$$

which is (4).

Proof of Proposition 1. The expected cash flow is $\mathbb{E}[X] = \bar{\Pi} - \theta(S - S^{\text{ref}}) - SV\bar{p}$, so

$$\mathbb{E}[X] - P^*(S) = \lambda bSV \qquad \text{and} \qquad \mathbb{E}[r^*(S)] = \frac{\lambda bSV}{P^*(S)},$$

which is (6). Since $dP^*/dS = -[\theta + V(\bar{p} + \lambda b)]$,

$$\frac{d\mathbb{E}[r^*]}{dS} = \frac{\lambda bV\big\{P^*(S) + S[\theta + V(\bar{p} + \lambda b)]\big\}}{P^*(S)^2} = \frac{\lambda bV(\bar{\Pi} + \theta S^{\text{ref}})}{P^*(S)^2}.$$

The term $\bar{\Pi} + \theta S^{\text{ref}} = P^*(0)$ is the value of a firm with no shortfall, which we take to be positive, so the derivative is positive for $\lambda b > 0$.

Proof of Proposition 2. From (5),

$$P^*(S^{\text{ref}}) - P^*(S) = (S - S^{\text{ref}})[\theta + V(\bar{p} + \lambda b)].$$

Let

$$N(S) \equiv \mathbb{E}[X] - P(S) = \lambda bSV + \mu\,(S^{\text{ref}} - S)[\theta + V(\bar{p} + \lambda b)],$$

so that $\mathbb{E}[r(S)] = N(S)/P(S)$, which is (7). Differentiating,

$$\frac{d\mathbb{E}[r]}{dS} = \frac{N'P - NP'}{P^2}, \qquad N' = \lambda b(1-\mu)V - \mu(\theta + \bar{p}V), \qquad P' = -(1-\mu)[\theta + V(\bar{p} + \lambda b)],$$

so that

$$\operatorname{sign}\left(\frac{d\mathbb{E}[r]}{dS}\right) = \operatorname{sign}\Big(N' + \mathbb{E}[r]\,(1-\mu)[\theta + V(\bar{p} + \lambda b)]\Big).$$

The second term is the expected return times the price sensitivity $(1-\mu)[\theta + V(\bar{p} + \lambda b)]$, which is of the same order as the terms in N'; when $|\mathbb{E}[r]| \ll 1$ it is negligible and the sign is that of N', which gives condition (8). Exactly, expected returns decrease in S if and only if

$$\mu(\theta + \bar{p}V) > \lambda b(1-\mu)V + \mathbb{E}[r]\,(1-\mu)[\theta + V(\bar{p} + \lambda b)].$$

Relation to Bolton et al. (2026). Their static model with financial constraints gives

$$\frac{\partial V}{\partial p} = g'(I^*)(A - C), \qquad g'(I^*) > 1,$$

where $g(\cdot)$ is their production function. The full-information value above with $\theta = 0$ is the unconstrained case; constraints scale the exposure SV by $g'(I^*)$ and leave the sign structure of (7) unchanged.

Appendix B EU ETS Institutional Background

The EU ETS is a cap-and-trade system that limits greenhouse gas emissions from power and heat generation, energy-intensive industry (including refining, iron and steel, cement and lime, glass, ceramics, pulp and paper, and chemicals) and, since 2012, aviation within the European Economic Area. Each regulated installation monitors its annual emissions, which an accredited verifier audits. The operator must report verified emissions for a calendar year by 31 March of the following year and surrender one allowance for each tonne of CO_2-equivalent. Until the 2023 compliance year, the surrender deadline was 30 April; it has since moved to 30 September. An operator that fails to surrender enough allowances pays an excess-emissions penalty of €100 per tonne, indexed to inflation since

2013, and must still surrender the missing allowances. Compliance is therefore binding at the installation level.

The system has developed in four phases. Phase 1 (2005–2007) was a pilot phase in which allowances were allocated generously and almost entirely for free; Ellerman and Buchner (2007) discuss the evidence on over-allocation in the first compliance year, and Anderson and Di Maria (2011) estimate that both over-allocation and abatement occurred over the pilot phase. Phase 1 allowances could not be banked into Phase 2, and their price collapsed to essentially zero by mid-2007 (Hintermann 2010). In Phase 2 (2008–2012), member states continued to set national allocation plans and to allocate most allowances for free, largely on the basis of historical emissions. The recession of 2008–2009 reduced output and emissions, and allowance prices fell.

Phase 3 (2013–2020) replaced national plans with an EU-wide cap and harmonized allocation rules. Auctioning became the default method of allocation, and electricity generators no longer received free allocation, except for transitional allocation under Article 10c of the ETS Directive in some lower-income member states. Industrial installations continued to receive free allocation, calculated from product benchmarks based on the most efficient installations in each sector and adjusted for sectors deemed exposed to carbon leakage. New installations and capacity extensions could receive allowances from the new entrants reserve. The Market Stability Reserve, operating since 2019, withdraws allowances from circulation when the surplus is large.

Phase 4 (2021–2030) tightened the cap and updated the benchmark values. Under the 2023 revision of the ETS Directive, free allocation to sectors covered by the carbon border adjustment mechanism is to be phased out gradually. Allowance prices rose from below €10 for most of 2012–2017 to above €80 for much of 2022 and 2023.

These rules have two consequences for the allocation shortfall. First, the shortfall of electricity generators rose abruptly in 2013, when their free allocation ended, and has since been mechanically close to one; for this reason the main sample excludes utilities. The

move to harmonized, benchmark-based allocation in 2013 also explains the jump in the shortfall of the remaining firms in Figure 1. Second, the shortfall of an industrial installation depends on how its emissions compare with a product benchmark, on the historical activity level to which the benchmark is applied and on its sector's carbon-leakage status. The first two reflect the installation's operations and the last its regulatory treatment, and all three change slowly, which is why the persistent component of the shortfall captures a slow-moving regulatory position.

The European Union Transaction Log records verified emissions, allocations and surrendered allowances for each installation and year, together with account holders and transactions. We use the relational version of the EUTL compiled by Abrell (2024), which links installations, accounts and account holders over time.

Appendix C EUTL–Compustat Matching Procedure

This appendix summarizes the construction of the installation-to-parent mapping. The full documentation, code and logs of every database query are available from the authors.

C.1 Sources

We combine four sources: the EU ETS data package of Abrell (2024) (2026 snapshot), which provides compliance data, installations, accounts, account holders and transactions; the 2024 vintage of the same data, which preserves links between installations and former operator accounts; Orbis (Bureau van Dijk) on WRDS, for identifiers, legal status, ownership links and equity ISINs (on the coverage of Orbis, see Bajgar et al. 2020); and Compustat Global and North America, for securities, prices, fundamentals and exchange rates. A legacy mapping based on legal entity identifiers (LEIs) serves as a cross-check and as a fallback.

C.2 Procedure

The procedure has five steps.

Installations and holders. We keep stationary installations (excluding aviation, maritime and ETS2 accounts). For each installation and year, we assign one account holder: the holder of the account open on 31 December; if none, the account with the largest overlap with the year; and otherwise the nearest spell. Every installation-year from 2005 to 2025 receives a holder. For years up to 2023, we use the 2024-vintage holder when it differs from the current one, to capture transfers of installations between holders.

Holders to Orbis. We match holders to Orbis entities through, in order, the LEI recorded in the EUTL, the LEI from the legacy mapping, national registration numbers (including constructed BvD identifiers and German register numbers with register-court agreement), Orbis identifiers recorded in the 2024 EUTL, exact domestic names, and 12 hand-verified anchor rows for 13 large holders. Every candidate is validated by name, after transliteration and removal of legal-form suffixes. About 99% of emissions over 2008–2023 are matched at this step. We checked the 100 largest holders, which account for 50% of emissions, by hand; 98 were correct, and both errors were fixed.

Control chain. From each matched entity, we follow the largest immediate shareholder upward as long as its stake exceeds 50%. Listed ancestors are nodes with listed or delisted status and an equity ISIN. When the chain reaches no listed node, we use Orbis's global ultimate owner at the 50% threshold.

Listed parents to Compustat. We link listed ancestors to Compustat through ISIN, SEDOL and CUSIP, including historical identifiers. For holders with no Orbis listed parent, we use the listed parent from the legacy LEI mapping.

Parent assignment by year. Orbis on WRDS provides current ownership links only. We therefore apply 34 dated overrides for known ownership changes that the snapshot cannot see (for example, Uniper to E.ON up to 2015, EDF Polska to EDF up to 2016 and to PGE afterward, and three cases identified in the literature comparison below). We then walk

up the ranked listed ancestors and take the first one that is priced in the year, skipping a node only if it was delisted before that year and stopping at a node that was not yet listed. The resulting "nearest" scheme is the baseline. The "consolidated" scheme also assigns emissions to upper listed ancestors within 35 documented control windows between listed subsidiaries and listed parents; it is used as a robustness check.

C.3 Confidence Flags and Validation

Each assignment carries a confidence flag. An assignment is marked low confidence if it relies on a fallback, skips a delisted node, involves a strongly renamed account, disagrees with the legacy mapping, uses a data-derived override, or reaches an inactive Orbis subsidiary. Low-confidence assignments account for 14.2% of the emissions assigned under the nearest scheme.

The mapping agrees with the legacy LEI-based mapping for 99.0% of the emissions that both assign. Between 53.1% and 59.1% of stationary emissions per year (2008–2025) reach a priced EU-listed parent, and there are 509 to 549 EU-listed parents per year. Monthly returns agree with the legacy panel (correlation 0.98 over 59,973 firm-months, median absolute difference zero), as do book-to-market ratios (correlation 0.997) and market capitalizations (median ratio 1.002). The shortfall has a Spearman correlation of 0.954 with the legacy panel's measure.

We also compare the sample with published EUTL studies. The sector mix is close to that of Fuchs (2025): at the installation level, combustion accounts for 70.1% of emissions, steel for 6.9%, cement for 8.3% and refining for 6.7%, against 64.6%, 7.5%, 6.5% and 6.0% in Fuchs's sample of operators. The ratio of allocation to verified emissions in the listed sample (0.72 over all phases; 0.99 in Phase 1 and 0.45 in Phase 4) is close to the values computed from Laurs et al. (2026) (0.70; 1.00 and 0.43). Listed parents cover 61% of EU ETS emissions in this mapping over 2005–2022, against about 74% in Laurs et al. (2026). Most of the gap is definitional: Laurs et al. attribute emissions to listed shareholders in

proportion to their direct and indirect stakes, including minority and joint-venture stakes, whereas this mapping requires a majority control chain to a priced listed parent. The largest unassigned holders are unlisted, such as state-owned utilities and privately held generators, or are 50/50 joint ventures.

C.4 Sensitivity to Mapping Quality

A diagnostic script recomputes the shortfall after removing installations with doubtful assignments, in the main sample without utilities. It uses a slightly different implementation of the sort from the one in the main text, and its baseline L−H spread is 4.9% per year (with a t-statistic of 2.30) rather than 5.4%. Excluding installations assigned through overrides, fallbacks, anchors or post-delisting skips, computing the shortfall from high-confidence installations only, or dropping firm-years in which more than half of emissions are low confidence leaves the spread between 4.3% and 5.4% per year, significant at the 5% level in every case. Dropping firm-years with any low-confidence emissions, which removes a quarter of firm-months, gives a spread of 5.1% per year that is significant at the 10% level. The persistent shortfall in the same diagnostic gives 5.3% per year, with a t-statistic of 2.38. Returns are annualized by multiplying monthly values by 12.

C.5 Known Limitations

Ownership links are a current snapshot, corrected only for documented changes. The pre-IPO years of later-listed state utilities are unassigned by construction. Joint ventures with 50/50 ownership are not assigned. Parents headquartered outside Europe are priced only on their European lines and are outside the baseline sample. About 1% of emissions belong to holders whose registration-number formats are not yet handled.

Appendix D Additional Tables

This appendix reports the additional tables referred to in the main text. Table D1 reports the main results for the full sample including utilities. Tables D2 and D3 report robustness tests for the shortfall level and the persistent shortfall.

Table D4 reports equal-weighted portfolios and Fama–MacBeth regressions. Neither detects the premium: the equal-weighted L−H spreads are positive but insignificant, and the Fama–MacBeth slopes on the shortfall are insignificant. Both methods give every firm the same weight, whereas Section 5.4 shows that the premium is earned among the large, widely followed firms to which institutional investors and analysts pay attention, and is absent among small firms and firms that no analyst follows. Equal weighting therefore dilutes the premium with firms whose carbon information few investors process. Table D5 shows that excluding microcaps does not change the value-weighted results.

Table D6 reports double sorts on ESG ratings and on verified emission intensity. Controlling for the ESG score or the environmental-pillar score, the average spread across terciles is positive but not statistically significant, and its five-factor alpha is significant only at the 10% level.

Table D7 splits the sample at the median of the economic materiality of the EU ETS (Section 5.1).

Table D1: Main Results in the Full Sample Including Utilities

This table reports value-weighted quintile L−H spreads and alphas (% per year) for the full sample, which includes utilities (SIC 4900–4999 or GICS sector 55): 5,847 firm-years for 562 firms, formation years 2009–2023. The signals and the construction are those of Tables 2 and 4; the persistent signals use the firm-years with both components (5,662). Returns and alphas are annualized by multiplying monthly values by 12. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024.

	L−H		CAPM		FF3		FF5		FF5+MOM		
Signal	Ret	t	α	t	α	t	α	t	α	t	Months
Shortfall level	5.03	2.22	5.87	2.48	5.19	2.05	5.43	2.13	4.57	1.72	180
Persistent shortfall	6.29	2.81	7.55	3.21	6.90	2.73	7.18	2.83	6.33	2.42	180
Persistent ⊥ transitory	7.16	2.93	8.07	3.10	7.48	2.82	8.09	3.05	7.71	2.74	180

Table D2: Robustness Tests

This table reports value-weighted quintile L−H spreads (% per year) on the shortfall level, May 2009–April 2024 (180 months). Returns and alphas are annualized by multiplying monthly values by 12. The weight cap limits any firm to 10% of its portfolio. The consolidated scheme also assigns emissions to upper listed parents within documented control windows. The NER/Art. 10c row excludes the new entrants reserve and Article 10c allocation from A. Truncation at −1 only creates ties within Q1 and therefore reproduces the baseline sort. The split by the economic materiality of the EU ETS is in Table D7. GB-dominant firm-years are those with most EU ETS emissions in the United Kingdom. Fama-MacBeth estimates are in Table D4. t-statistics use Newey-West standard errors with 6 lags.

	L−H		CAPM		FF5		FF5+MOM		
	Ret	t	α	t	α	t	α	t	Months
Panel A: Weighting and portfolio design									
Baseline: VW quintiles	5.40	2.62	6.19	2.77	6.86	2.65	6.15	2.17	180
Equal-weighted quintiles	1.56	0.86	2.07	1.12	2.12	1.09	1.51	0.81	180
VW, 10% weight cap	4.00	2.02	4.01	1.85	4.62	1.91	3.78	1.50	180
Deciles (D1−D10)	5.58	1.86	5.80	1.83	5.24	1.72	3.41	1.06	180
Terciles (T1−T3)	3.98	2.58	3.70	2.26	4.37	2.60	3.79	2.08	180
Panel B: Data and signal definition									
Consolidated ownership scheme	5.16	2.48	6.00	2.65	6.72	2.60	6.07	2.15	180
Allocation excl. NER and Art. 10c	3.22	1.49	4.03	1.92	3.89	1.61	3.19	1.25	180
Truncated at −1	5.40	2.62	6.19	2.77	6.86	2.65	6.15	2.17	180
Panel C: Sample									
Excluding GB-dominant firm-years	5.59	2.80	5.79	2.70	5.82	2.50	5.68	2.28	180
SIC2-demeaned signal	3.47	1.75	3.14	1.65	3.27	1.50	3.16	1.38	180
Quintiles within SIC2	3.84	1.51	4.13	1.68	3.86	1.51	3.65	1.34	180

Table D3: Robustness of the Persistent Shortfall

This table reports value-weighted quintile L−H spreads (% per year) on the persistent shortfall (Panel A) and on the persistent shortfall orthogonalized to the transitory shortfall (Panel B), May 2009–April 2024. The within-industry rows demean the signal within SIC2 industry and formation year. Panel C drops one firm at a time from the persistent sort (478 firms) and reports the range of the resulting L−H spreads and t-statistics; Share $t < 2$ is the fraction of drops that leave a t-statistic below 2. Returns and alphas are annualized by multiplying monthly values by 12. t-statistics use Newey-West standard errors with 6 lags.

	L−H		CAPM		FF5		FF5+MOM		
	Ret	t	α	t	α	t	α	t	Months
Panel A: Persistent shortfall									
Baseline	5.76	2.68	6.78	2.91	7.57	3.21	6.09	2.48	180
VW, 10% weight cap	4.40	2.30	4.49	2.14	5.16	2.49	3.73	1.75	180
Equal-weighted	2.48	1.60	2.86	1.75	2.59	1.45	1.58	0.95	180
Within industry (SIC2-demeaned)	4.55	2.02	4.46	1.92	4.03	1.51	3.84	1.44	180
Panel B: Persistent shortfall ⊥ transitory									
Baseline	6.74	3.43	7.70	3.50	7.75	3.38	6.91	2.94	180
VW, 10% weight cap	5.39	2.65	5.72	2.58	5.97	2.79	5.34	2.51	180
Equal-weighted	2.01	1.23	2.48	1.39	2.24	1.20	1.78	0.94	180
Within industry (SIC2-demeaned)	3.45	1.62	3.46	1.57	3.38	1.41	2.85	1.20	180

	Full sample		L−H, one firm dropped			t, one firm dropped			Share
Firms	L−H	t	Min	Median	Max	Min	Median	Max	$t < 2$
Panel C: Leave-one-firm-out, persistent shortfall									
478	5.76	2.68	3.54	5.95	6.88	1.84	2.73	3.05	0.002

Table D4: Equal-Weighted Portfolios and Fama-MacBeth Regressions

Panel A compares value-weighted (VW) and equal-weighted (EW) L−H quintile spreads (% per year; monthly values multiplied by 12) for the shortfall level, the persistent shortfall and the persistent shortfall orthogonalized to the transitory shortfall, May 2009–April 2024 unless stated; Q1 and Q5 are the average excess returns of the extreme quintiles. Panel B reports Fama and MacBeth (1973) slopes (%/month per standard deviation) of monthly returns on the signal standardized within formation year; controls are log market capitalization, B/M, operating profitability, asset growth, leverage and $\text{Mom}_{12\text{-}2}$. A *negative* slope corresponds to a positive L−H spread. N is the average number of firms per cross-section. t-statistics use Newey-West standard errors with 6 lags.

	L−H		CAPM		FF5		FF5+MOM		Q1	Q5
	Ret	t	α	t	α	t	α	t	Ret	Ret
Panel A: Value- versus equal-weighted L−H spreads										
Shortfall level, VW (baseline)	5.40	2.62	6.19	2.77	6.86	2.65	6.15	2.17	11.44	6.04
Shortfall level, EW	1.56	0.86	2.07	1.12	2.12	1.09	1.51	0.81	12.95	11.38
Persistent shortfall, VW	5.76	2.68	6.78	2.91	7.57	3.21	6.09	2.48	13.57	7.82
Persistent shortfall, EW	2.48	1.60	2.86	1.75	2.59	1.45	1.58	0.95	14.39	11.91
Persistent ⊥ transitory, VW	6.74	3.43	7.70	3.50	7.75	3.38	6.91	2.94	13.66	6.92
Persistent ⊥ transitory, EW	2.01	1.23	2.48	1.39	2.24	1.20	1.78	0.94	13.57	11.57

Signal	Controls	Sample	Slope	t	Months	N
Panel B: Fama-MacBeth regressions						
Shortfall	No	2009:05–2024:04	0.02	0.48	180	333
Shortfall	Yes	2009:05–2024:04	0.07	1.42	180	332
Persistent shortfall	Yes	2009:05–2024:04	0.07	1.31	180	322

Table D5: Excluding Microcaps

This table reports value-weighted (VW) and equal-weighted (EW) quintile L−H spreads and five-factor alphas (% per year; monthly values multiplied by 12) for all firms (Panel A) and after excluding firm-years whose end-of-April market capitalization is below the 20th percentile of all European-exchange primary common stocks in Compustat Global in that April (Panel B; breakpoints between USD 7.7 and 16.0 million over formation years 2009–2023). The signals are those of Table 4. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024.

	Value-weighted				Equal-weighted			
Signal	L−H	t	FF5 α	t	L−H	t	FF5 α	t
Panel A: All firms								
Shortfall level	5.40	2.62	6.86	2.65	1.56	0.86	2.12	1.09
Persistent shortfall	5.76	2.68	7.57	3.21	2.48	1.60	2.59	1.45
Persistent ⊥ transitory	6.74	3.43	7.75	3.38	2.01	1.23	2.24	1.20
Panel B: Excluding microcaps (below the 20th percentile of European market capitalization)								
Shortfall level	5.36	2.60	6.80	2.59	−0.10	−0.06	0.13	0.07
Persistent shortfall	5.61	2.57	7.48	3.16	1.63	1.12	1.45	0.90
Persistent ⊥ transitory	6.77	3.40	7.98	3.53	0.59	0.36	0.65	0.35

Table D6: Double Sorts Controlling for ESG Ratings and Verified Emission Intensity

This table repeats the double sort of Table 5 with further controls. Panel A sorts the shortfall level within terciles of the LSEG ESG score and environmental-pillar (E-pillar) score for the latest fiscal year ending by December of year $t-1$ (about 180 firms per formation year). LSEG revises its historical ESG scores (Berg et al. 2020); we use the current vintage, so the scores may contain information that was not available at formation. Panel B uses terciles of verified EU ETS emission intensity (verified emissions of compliance year $t-1$ over sales, tonnes per USD million); the persistent shortfall uses the sample of Table 4. Average is the monthly mean of the three tercile L−H returns; FF5 α is its five-factor alpha. Unconditional is the one-way value-weighted L−H spread in the same sample. Returns and alphas are annualized by multiplying monthly values by 12. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024.

	T1 (Low)		T2		T3 (High)		Average		Average		Unconditional	
Control / signal	L−H	t	L−H	t	L−H	t	L−H	t	FF5 α	t	L−H	t
Panel A: Shortfall level, controlling for ESG ratings												
ESG score	−1.84	−0.40	8.32	2.67	2.36	1.09	2.95	1.48	3.40	1.69	6.00	3.11
E-pillar score	−4.89	−1.22	6.74	2.46	5.00	1.91	2.28	1.34	3.05	1.92	6.00	3.11
Panel B: Control is verified EU ETS emission intensity												
Shortfall level	4.79	1.68	8.15	1.65	0.54	0.10	4.49	2.08	2.96	1.38	5.43	2.65
Persistent shortfall	5.77	2.14	3.74	0.82	0.51	0.10	3.34	1.16	3.19	1.08	5.98	2.75

Table D7: The Shortfall Spread by Materiality of the EU ETS

At the end of April of each year, firms are split at the within-year median of the materiality of the EU ETS, independently of the signal, and value-weighted terciles (Panels A, B and D) or quintiles (Panel C, at least three firms per quintile) of the signal are formed within each half. In Panels A–C, materiality is verified emissions of compliance year $t-1$ over sales in U.S. dollars of the latest fiscal year (tonnes per USD million). In Panel D, it is verified emissions over LSEG Scope 1 emissions of fiscal year $t-1$, capped at one (about 146 firms per formation year). Panel B demeans the signal within SIC2 industry and formation year. L−H is the lowest- minus the highest-shortfall portfolio within the half (% per year); FF5 α is its Fama–French five-factor alpha; High − Low is the difference between the L−H returns of the high- and the low-materiality halves. The persistent signals use the sample of Table 4. The average number of firms per portfolio leg is 53 to 55 in Panels A and B, 32 to 33 in Panel C and 23 to 24 in Panel D. Returns and alphas are annualized by multiplying monthly values by 12. t-statistics use Newey-West standard errors with 6 lags. Sample: May 2009–April 2024 (180 months).

	Low materiality				High materiality				High − Low	
Signal	L−H	t	FF5 α	t	L−H	t	FF5 α	t	L−H	t
Panel A: Materiality = verified emissions / sales; signal terciles										
Shortfall level	4.77	3.12	4.75	2.90	−1.39	−0.37	−3.35	−1.03	−6.15	−1.47
Persistent shortfall	3.82	2.45	4.20	2.48	−0.89	−0.22	−3.43	−0.97	−4.70	−1.05
Persistent ⊥ transitory	4.09	2.59	3.90	2.41	2.04	0.49	0.62	0.16	−2.06	−0.46
Panel B: As Panel A, signal demeaned within SIC2 industry										
Shortfall level	2.67	1.44	2.49	1.40	−2.21	−0.83	−2.20	−0.79	−4.88	−1.51
Persistent shortfall	5.07	2.67	4.93	2.39	0.41	0.17	0.93	0.36	−4.66	−1.55
Persistent ⊥ transitory	5.24	2.39	4.80	1.84	−1.91	−0.74	−1.27	−0.49	−7.15	−2.26
Panel C: As Panel A, signal quintiles										
Shortfall level	4.01	1.68	4.80	1.89	4.11	0.90	1.63	0.39	0.09	0.02
Persistent shortfall	4.77	1.84	5.35	1.73	1.05	0.25	−0.45	−0.11	−3.72	−0.78
Persistent ⊥ transitory	8.80	3.44	10.14	3.30	4.75	0.94	3.91	0.92	−4.05	−0.76
Panel D: Materiality = verified emissions / LSEG Scope 1; signal terciles										
Shortfall level	6.15	2.97	5.45	2.55	−1.87	−0.65	−0.58	−0.22	−8.01	−2.07
Persistent shortfall	3.66	2.07	3.38	1.61	0.12	0.04	0.66	0.24	−3.55	−1.03
Persistent ⊥ transitory	2.93	1.42	2.20	0.92	0.86	0.29	1.44	0.46	−2.08	−0.56

Internet Appendix

Beyond Emissions: Regulatory Carbon Positions and Stock Returns

Internet Appendix: Secondary Tables and Figures

This Internet Appendix collects secondary tables and figures: the composition of the sort sample (Table IA1), the largest holdings of the extreme quintiles (Table IA2), the quintiles of the persistent-shortfall sort (Table IA3), risk characteristics by quintile (Table IA4) and the carbon cost burden by phase (Table IA5); and the leave-one-firm-out distribution of the persistent spread, rolling L−H spreads, allowance-price returns next to the L−H return, and the cumulative L−H return with and without utilities (Figures IA1 to IA4).

Table IA1: Sample Composition by Country and Sector

This table reports averages over formation years 2009–2023 of the sort sample by the country with most of the firm's EU ETS emissions (top ten by market capitalization) and by the firm's dominant EU ETS activity.

Country	Firms per year	Share of firms	Share of mkt cap	Activity	Firms per year	Share of firms	Share of mkt cap
DE	58.3	0.175	0.248	Combustion	188.8	0.561	0.722
FR	24.7	0.074	0.119	Refining	15.8	0.047	0.105
GB	29.0	0.084	0.108	Chemicals	10.6	0.032	0.043
ES	23.6	0.070	0.088	Cement and lime	18.8	0.056	0.034
IT	20.3	0.061	0.081	Iron and steel	23.9	0.071	0.028
NL	10.7	0.032	0.058	Pulp and paper	32.5	0.097	0.026
BE	14.6	0.044	0.052	Non-ferrous metals	11.0	0.033	0.025
AT	11.2	0.033	0.052	Glass, ceramics	32.1	0.095	0.014
IE	5.7	0.017	0.047	Other	1.2	0.004	0.003
PL	33.1	0.099	0.040	Coke	1.3	0.004	0.000

Table IA2: Largest Holdings of the Extreme Quintiles

This table lists the ten largest holdings of the lowest- (Panel A) and highest-shortfall (Panel B) quintiles in the level sort, ranked by end-of-April value weight averaged over all 15 formation years (2009–2023). Years is the number of formation years in which the firm is in the quintile. ETS country is the country with most of the firm's EU ETS emissions. Median shortfall is the firm's median raw shortfall.

Firm	SIC2	ETS country	Years	Avg. weight	Median shortfall
Panel A: Q1 (low shortfall)					
Roche Holding	28	DE	7	0.130	−1.24
Novo Nordisk	28	DK	8	0.115	−295
Novartis	28	AT	4	0.070	−0.11
ArcelorMittal	33	FR	10	0.039	−0.15
Volvo	37	SE	10	0.031	−0.58
Unilever	28	IT	4	0.030	−1.96
Reckitt Benckiser	28	ES	4	0.023	−1.09
Svenska Cellulosa	24	SE	10	0.020	−1.83
UPM-Kymmene	26	DE	11	0.019	−0.38
UCB	28	BE	8	0.019	−0.33
Panel B: Q5 (high shortfall)					
Vodafone	48	IE	15	0.109	1.00
British American Tobacco	21	PL	9	0.060	0.69
Equinor	29	NO	9	0.058	1.00
Sanofi	28	FR	5	0.047	0.13
Nestlé	20	ES	3	0.042	0.77
Eni	13	IT	8	0.041	0.61
Volkswagen	37	DE	15	0.039	0.72
Anheuser-Busch InBev	20	BE	5	0.034	0.69
ASTRAZENECA PLC	28	IE	4	0.029	0.49
Reckitt Benckiser	28	ES	5	0.025	0.72

Table IA3: Quintiles Sorted on the Persistent Shortfall

This table reports average excess returns and alphas (% per year) of value-weighted quintiles sorted on the persistent shortfall (three-year mean), May 2009–April 2024. Returns and alphas are annualized by multiplying monthly values by 12. t-statistics use Newey-West standard errors with 6 lags.

	Mean		CAPM		FF3		FF5		FF5+MOM		
	Ret	t	α	t	α	t	α	t	α	t	N
Q1 (Low)	13.57	3.44	6.93	3.50	6.77	3.26	6.76	3.16	5.32	2.27	64
Q2	9.67	2.13	0.75	0.54	0.87	0.63	−1.04	−0.67	−0.34	−0.18	65
Q3	9.88	2.23	2.08	1.05	2.56	1.49	0.20	0.11	0.59	0.29	64
Q4	9.81	2.55	1.99	1.24	2.17	1.45	1.45	0.88	1.43	0.87	64
Q5 (High)	7.82	2.02	0.16	0.10	0.38	0.25	−0.81	−0.60	−0.77	−0.49	65
L−H	5.76	2.68	6.78	2.91	6.38	2.56	7.57	3.21	6.09	2.48	

Table IA4: Risk Characteristics by Quintile

Panel A uses monthly portfolio returns against the European market factor, May 2009–April 2024. Volatility is annualized (%). Downside (upside) beta uses months in which the market excess return is below (above) its sample mean (Ang et al. 2006). Maximum drawdown is the largest peak-to-trough decline of the compounded return (%); the worst month is in %. L−H is treated as a self-financing return. Panel B uses local-currency daily returns from May of the year before formation to April of the formation year (at least 120 days); beta and downside beta are estimated against the value-weighted daily return of the sample universe (downside: days with the market below its window mean, at least 50 days), and idiosyncratic volatility is the residual volatility of this market model, because daily European factors are not available. VW is the end-of-April value-weighted mean within the quintile, Med the median, both averaged over formation years; t(Q1−Q5) is computed across the 15 formation years; Cov. is the share of firm-years with a daily estimate.

	Volatility (% p.a.)	Market beta	t	Downside beta	Upside beta	Max. drawdown	Worst month	Skewness
Panel A: Portfolio level								
Q1 (Low)	16.45	0.82	17.34	0.87	0.86	−22.52	−14.27	−0.26
Q2	19.56	1.05	26.23	1.15	1.09	−33.82	−19.97	−0.11
Q3	17.88	0.90	22.69	0.88	0.96	−24.78	−13.50	0.16
Q4	16.75	0.88	37.74	0.89	0.94	−21.93	−12.76	0.00
Q5 (High)	17.39	0.91	30.26	0.99	0.94	−36.59	−14.91	−0.12
L−H	9.08	−0.09	−1.53	−0.12	−0.09	−15.53	−7.15	0.40

	Total vol.		Idiosyncratic vol.		Beta		Downside beta		
	VW	Med	VW	Med	VW	Med	VW	Med	Cov.
Panel B: Firm level, daily returns in the year before formation									
Q1 (Low)	27.39	34.10	22.16	29.06	0.94	0.75	0.97	0.85	0.95
Q2	27.96	33.11	20.99	28.09	1.10	0.90	1.12	1.00	0.96
Q3	27.26	31.81	20.38	27.34	1.07	0.79	1.09	0.87	0.97
Q4	25.14	30.34	18.69	26.56	0.98	0.74	0.95	0.79	0.97
Q5 (High)	26.33	30.98	20.48	27.12	0.96	0.76	0.95	0.83	0.95
Q1−Q5	1.06	3.12	1.68	1.95	−0.01	−0.01	0.02	0.02	
t(Q1−Q5)	1.09	4.07	2.02	2.59	−0.36	−0.36	0.38	0.62	

Table IA5: Aggregate Carbon Cost Burden by Phase and Quintile

This table reports the aggregate net carbon cost of each quintile (total cost over total sales or positive operating income, in %), averaged over the formation years of each phase (P2: 2009–2012; P3: 2013–2020; P4: 2021–2023). Net cost counts only shortfalls.

	Burden / sales					Net cost / sales					Burden / operating income				
Phase	Q1	Q2	Q3	Q4	Q5	Q1	Q2	Q3	Q4	Q5	Q1	Q2	Q3	Q4	Q5
P2	−0.105	−0.119	−0.049	−0.016	0.039	0.000	0.000	0.000	0.000	0.039	−1.020	−0.990	−0.464	−0.157	0.302
P3	−0.048	−0.005	0.014	0.031	0.035	0.000	0.007	0.017	0.032	0.035	−0.463	−0.017	0.157	0.260	0.308
P4	−0.099	0.174	0.158	0.306	0.057	0.006	0.175	0.158	0.306	0.057	−0.659	1.562	1.414	1.857	0.380

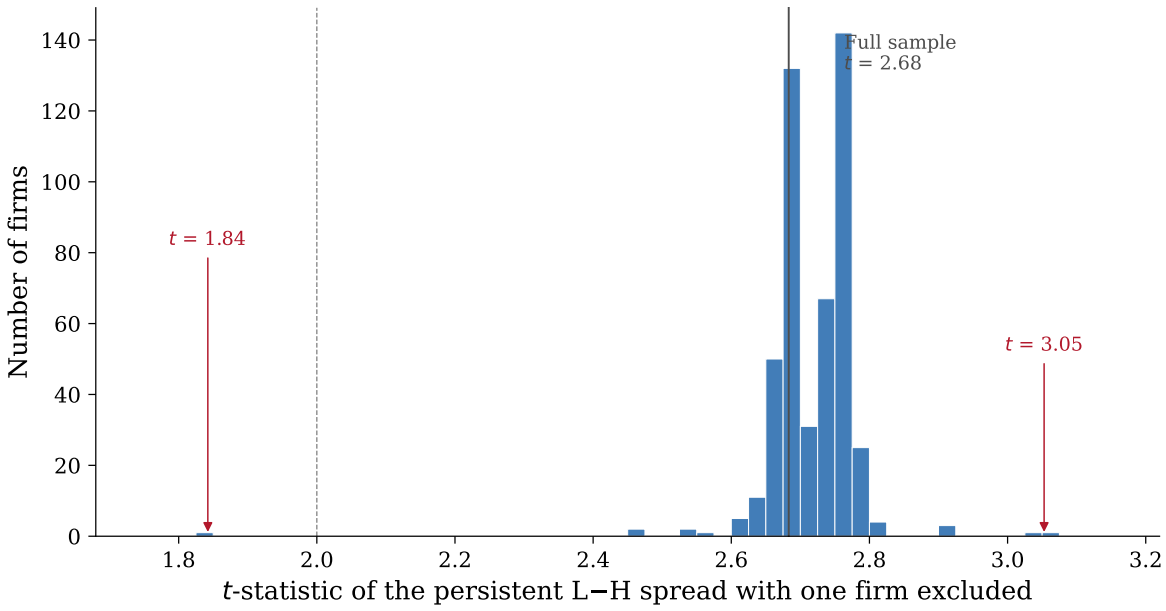


Figure IA1: **Leave-One-Firm-Out Distribution of the Persistent Spread.** This figure plots the distribution of the Newey-West t-statistic of the persistent-shortfall L−H spread when each of the 478 firms is excluded in turn. The solid line marks the full-sample t-statistic and the dashed line $t = 2$; arrows mark the two exclusions that change the t-statistic most.

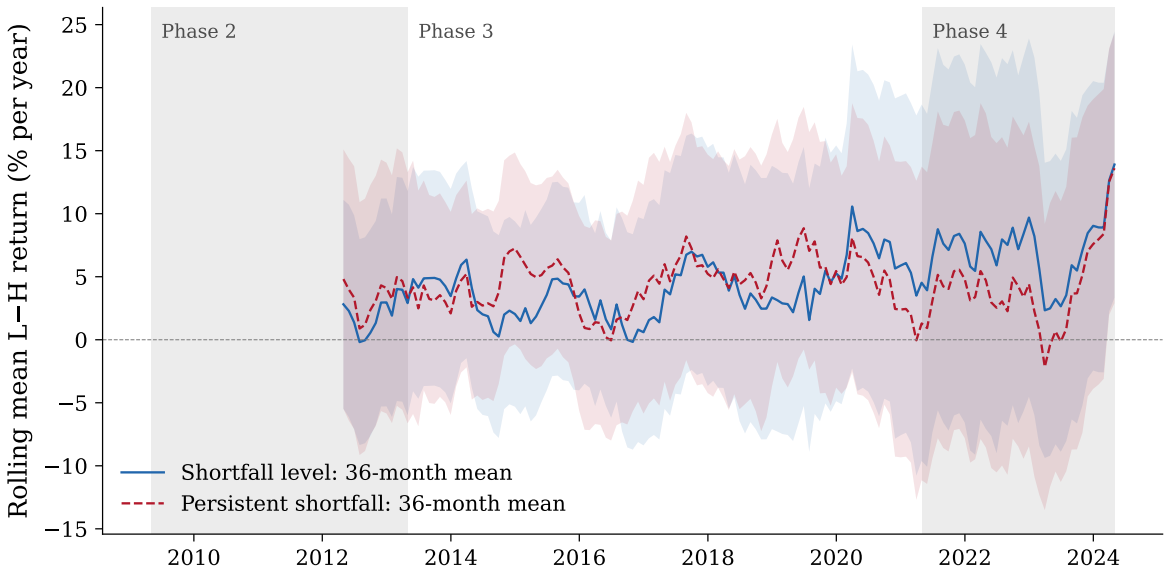


Figure IA2: **Rolling 36-Month L−H Spread.** This figure plots the 36-month rolling mean of the monthly value-weighted L−H return for the shortfall level (solid) and the persistent shortfall (dashed). Returns are annualized by multiplying monthly values by 12. Bands are ±1.96 rolling standard errors (iid). Background shading indicates the holding periods of Phase 2 and Phase 4.

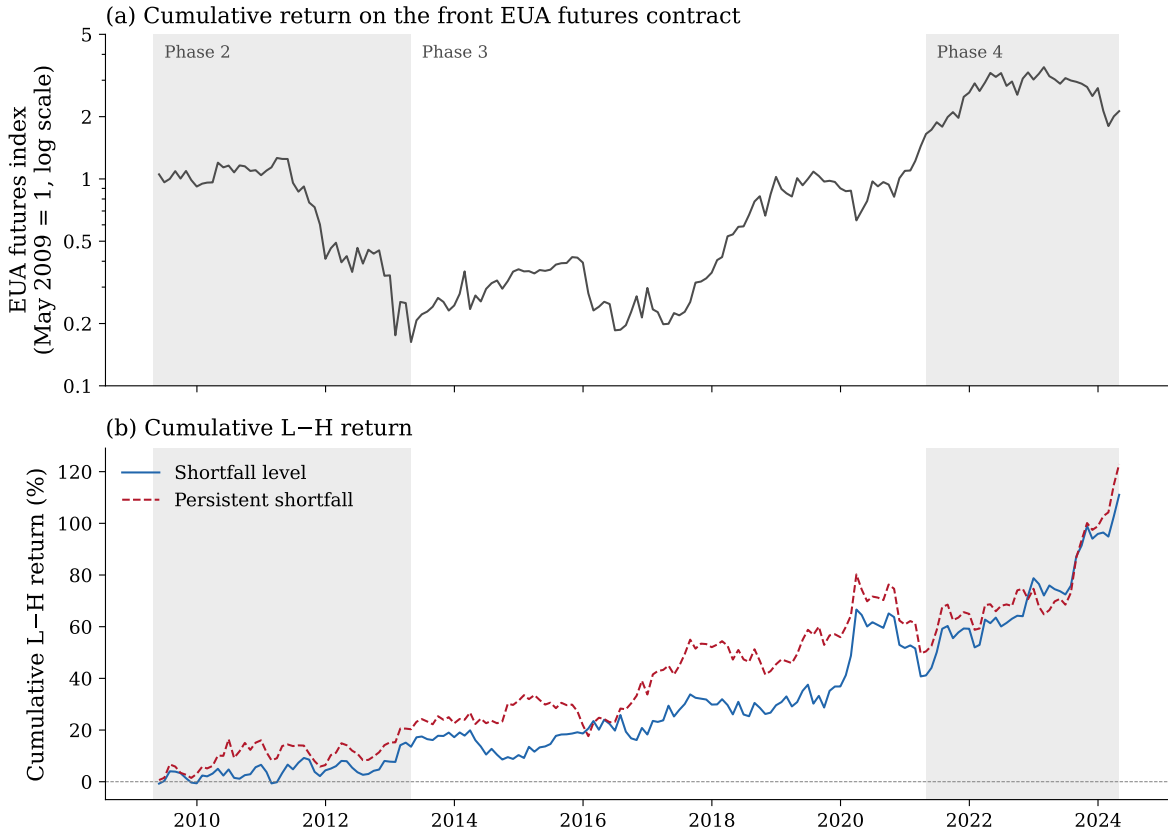


Figure IA3: **EUA Futures and the L−H Spread.** Panel (a) plots the cumulative return index (log scale) of the front EUA futures contract, compounded from within-contract daily settlement returns with roll days dropped. Panel (b) plots the cumulative value-weighted L−H return for the shortfall level and the persistent shortfall. Background shading indicates the holding periods of Phase 2 and Phase 4.

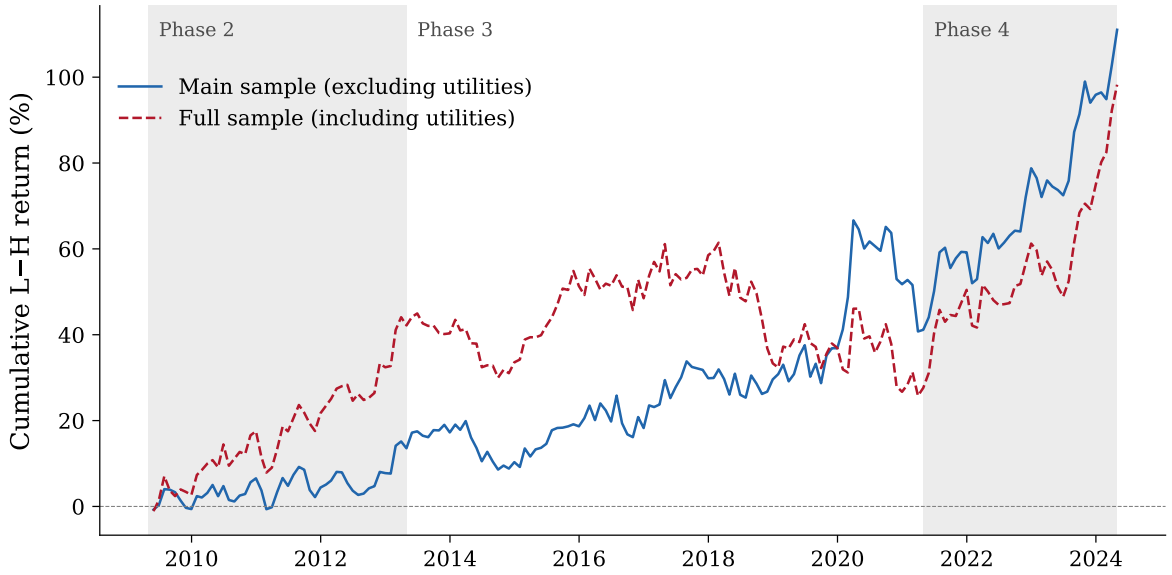


Figure IA4: **Cumulative L−H Returns with and without Utilities.** This figure plots the cumulative (compounded) return of the value-weighted L−H quintile portfolio sorted on the shortfall level in the main sample, which excludes utilities (SIC 4900–4999 or GICS sector 55; solid), and in the full sample including utilities (dashed), from May 2009 to April 2024. Background shading indicates the holding periods of EU ETS Phase 2 (formation years 2009–2012) and Phase 4 (2021–2023). The full-sample results are in Appendix Table D1.