



EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Olivier De Jonghe, Klaas Mulier,
Glenn Schepens, Leonard Stimpfle

Climate regulation, firm emissions,
and green takeovers

No 3253

Abstract

We show that an unexpected tightening of the EU Emissions Trading System led high-emission-intensity firms to cut emissions relative to low-intensity peers within the same industry, without reducing output, thereby improving emission efficiency. Effects are stronger for power producers than for manufacturing firms. Examining mergers and acquisitions (M&As), we find that high-intensity manufacturing firms acquire more green targets after the tightening than low-intensity firms, with no change in the overall number of acquisitions, indicating a shift in focus rather than activity. Finally, we show that these green M&As contributed to the observed emission reductions over the study period.

Keywords: climate regulation, emission trading, firm behavior, M&A

JEL classification codes: D22, G34, G38, Q53, Q54

NON-TECHNICAL SUMMARY

Since 2005 the EU Emissions Trading System (EU ETS) has capped emissions and required regulated installations to surrender allowances for each tonne of CO₂ emitted. For years, allowance prices were low and the system provided only weak incentives to cut emissions. That changed after 2017, when a series of regulatory reforms credibly tightened the cap and triggered a sharp, sustained rise in carbon prices. We use this period as a policy shock to examine whether and how firms adjusted their emissions.

We build a firm-level dataset covering 2,207 EU ETS firms between 2013 and 2023, combining verified emissions from the EU ETS registry with financial statements and merger-and-acquisition (M&A) data. The firms in our sample own a total of 5,314 regulated installations and collectively account for approximately 90% of total verified emissions in the EU ETS. To isolate the impact of the tightening, we compare firms within the same 4-digit industry that were more exposed to higher carbon prices (i.e., those with higher pre-2017 emission intensity) to otherwise similar firms that were less exposed. This design allows us to control for industry demand and other shared shocks while focusing on differences in carbon-price exposure.

The first result is a clear reduction in emissions among more exposed firms. After 2017, high-exposure firms cut emissions by about 13% relative to low-exposure firms in the same industry. The response is heterogeneous across sectors: power producers reduce emissions by roughly 32%, while manufacturing firms reduce emissions by about 5%. These effects build over time and strengthen further after the 2020 regulatory tightening, consistent with the continued rise in carbon prices.

The second result is that these emission cuts are achieved through efficiency improvements

rather than lower output. Operating revenue does not decline for more exposed firms after the tightening. Instead, emissions per unit of revenue fall by about 8.5% on average across all firms, with power producers improving efficiency by approximately 21% and manufacturing firms by approximately 5%. The dynamics, however, differ across sectors. For power producers, efficiency gains strengthen throughout the sample period. For manufacturing firms, efficiency improvements build until around 2020 but then weaken, and are no longer statistically significant by 2021. On average over the post-tightening period, the efficiency gain for manufacturing is meaningful, but the picture is less clear by the end of the sample. This indicates that higher carbon prices induced cleaner production rather than simple downsizing, though the persistence of this effect in manufacturing remains uncertain.

The third result highlights one of the channels at work: the market for corporate control. Overall M&A activity does not increase for more exposed firms, but their acquisition focus shifts toward greener targets. This change is concentrated in manufacturing firms, consistent with the view that fuel switching and other short-run abatement options are more readily available in the power sector. Green acquisitions are economically meaningful: firms reduce emissions by roughly 11% after completing a green deal, and excluding firms that ever do a green acquisition reduces the estimated emission response by about one-sixth in both the power and manufacturing sectors. These patterns suggest that acquisitions of cleaner firms and technologies are part of how manufacturing firms decarbonize in response to higher carbon prices. A further consequence is that carbon pricing generates positive spillovers beyond the directly regulated perimeter, by increasing the attractiveness of unregulated green firms as acquisition targets.

We also find no evidence that firms respond by shifting acquisitions outside the EU ETS area, which would indicate carbon leakage through the M&A channel. Financial constraints do not materially weaken the green acquisition response, likely reflecting the unusually ac-

commodative financing conditions during 2013–2023.

These findings are relevant for policy for three reasons. First, they show that carbon pricing can deliver real emission reductions when the price signal is strong and credible, even without reducing output. Second, they reveal that in hard-to-abate manufacturing sectors, acquisitions of cleaner firms and technologies are a practical mechanism for adjustment, potentially complementing other strategies such as in-house innovation. Third, the absence of increased cross-border acquisition leakage suggests that tightening the EU ETS can reduce emissions without simply moving activity elsewhere; our results also provide a pre-Carbon Border Adjustment Mechanism (CBAM) benchmark for assessing how future policy changes may affect the geography of green investment. Together, the results support the view that credible carbon pricing, combined with a functioning market for corporate control, can accelerate the transition to lower-emission production.

1. INTRODUCTION

Reducing carbon emissions through climate regulation is central to combating global warming and transitioning to a low-carbon economy. Since 2005, Emissions Trading Systems (ETS) have become a cornerstone of climate policy worldwide. An ETS creates a market for tradable emission allowances with an annual, often decreasing, limit on total emissions, thereby setting a uniform carbon price for regulated emitters. The number of active systems has evolved from just 1 in 2005 (the European Union ETS) to 36 in 2024, increasing the share of worldwide greenhouse gas emissions covered from around 4.5% to more than 18% (Figure OA.1).

The effectiveness of an ETS depends critically on how regulated firms adjust their behavior in response to carbon pricing. By placing an explicit cost on emissions, an ETS creates

financial incentives that can induce firms to mitigate carbon costs in different, mutually non-exclusive ways. For instance, firms could divest or downsize (Berg, Ma, and Streitz 2024; Boeckx, Struyfs, and Torsin 2025), they could decide to move production to non-regulated areas (Hanna 2010), or they could try to implement greener production technologies either by increasing in-house innovation (Calel and Dechezleprêtre 2016) or by acquiring firms with such technologies. Understanding firms' responses is essential for evaluating the actual impact of carbon pricing on climate outcomes.

This paper provides novel firm-level evidence on how firms respond to a series of regulatory tightening in the European Union Emission Trading System (EU ETS). While earlier studies focus on firm behavior in the initial phases of the EU ETS (Martin, Muûls, and Wagner 2016; Calel and Dechezleprêtre 2016; Dechezleprêtre, Nachtigall, and Venmans 2023; Colmer et al. 2025), the system has faced significant challenges and has undergone material changes in recent years which are, to date, understudied (Sitarz et al. 2024; Bruninx, Ovaere, and Delarue 2020). These changes started in 2017 with the announcement of an unexpected reform of the Market Stability Reserve, followed by the Climate Law in 2020, implemented with the Fit-for-55 package in 2021, which caused an unprecedented rally in carbon prices (Figure 1).

We exploit these post-2017 regulatory tightenings as a quasi-exogenous shock to the carbon price. Using a difference-in-differences strategy, we compare firms highly exposed to the carbon price shock to less exposed peer firms in the same 4-digit NACE industry over an eleven year period from 2013 to 2023. This setup allows us to control for product-market conditions while capturing variation in regulatory exposure. We define high-exposed firms as those with an average emission intensity during the pre-period (2013-2016) above the industry median. The higher a firm's emission intensity, the more emission allowances it has to surrender per unit of production and the more exposed it is to a marginal increase in

carbon prices.

Our analysis builds on a firm-year panel that combines installation-year level data on emissions from the EU ETS Union Registry¹ with firm-year level data on financial statements from Bureau van Dijk's Orbis database and firm-year level data on M&A deals from Bureau van Dijk's ZEPHYR database. Our sample consists of 24,202 firm-years from 2,207 firms over the period 2013–2023. These firms own a total of 5,314 regulated installations, covering around 90% of all regulated emissions in the EU ETS.

We find that high-exposed firms reduce emissions relative to their low-exposed peers by 13% on average over the post-regulatory tightening period (2017-2023).² These reductions translate into an 8.5% improvement in emissions per unit of operating revenue, suggesting that the emission reduction is not directly proportional to a reduction in output, and hence signals meaningful efficiency gains. The response is markedly heterogeneous. High-exposed firms in the power sector reduce emissions by 32% on average relative to their low-exposed peers, while high-exposed manufacturing firms show more modest relative reductions of about 5%. Importantly, event-study regression analyses show parallel trends in emissions and emission efficiency before 2017.

To investigate one of the channels through which firms reduce emissions, we examine changes in acquisition behavior. Motivated by the view that mergers and acquisitions (M&As) can serve as a least-cost adjustment mechanism in response to economic shocks (Mitchell and Mulherin 1996; Harford 2005), we test whether high-exposed firms shift toward acquiring “green” targets after 2017. Leveraging recent advancements in natural language processing, we employ the pretrained large language model ClimateBERT (Webersinke et al. 2022) to classify the environmental orientation of over 2,800 M&As announced by firms regulated by the EU ETS between 2013 and 2023.

Our analysis shows that high-exposed firms become significantly more likely to acquire green targets after the regulatory tightening relative to low-exposed firms. However, we do not find that the overall likelihood to do an acquisition increases for high-exposed firms relative to low-exposed firms. The increased acquisition of green targets therefore seems to reflect a change in M&A focus rather than an increase in M&A activity of high-exposed firms. Furthermore, the overall shift in acquisition behavior is almost entirely driven by high-exposed manufacturing firms with no statistically significant effect on power producers.

Next, we look directly and indirectly at the importance of these green M&As for the relative reduction in emissions by high-exposed firms. Regarding the direct evidence, within-firm analyses show that emissions significantly decline after a green deal is completed, where the estimates suggest an order of magnitude of around 10%. Regarding the indirect evidence, we find that excluding the firms that do a green deal at any point during our sample period weakens the relative emission reduction by almost 5 percentage points for high-exposed power firms (from 32.4% to 27.5%) and approximately 1 percentage point for high-exposed manufacturing firms (from 5.1% to 4.2%). In both the power and manufacturing industries, this represents around 16% of the observed emission reduction and is suggestive of the importance of the market for corporate control in firms' response to the regulatory tightening.

As financial conditions may matter for firms' capacity to engage in M&As, we analyse whether our findings differ for high-exposed, financially constrained firms than for high-exposed, financially unconstrained firms. We use ex-ante low cash holdings, high leverage, or small size as proxies for financial constraints. We find no evidence that financial constraints matter for our results using any of the measures. We interpret this null finding through the lens of the macroeconomic environment during our sample period (2013–2023). This period was characterized by exceptionally accommodative monetary policy, including negative policy rates, large-scale asset purchases, and ample market liquidity. Under these

conditions, even cash-strapped, highly levered, or small firms could likely secure external finance to fund their acquisitions.

We perform several robustness checks. For instance, we show that our results are not driven by our choice to define high-exposed firms as those with an above industry-median emission intensity. We show that the results hold when using the continuous emission intensity measure, and that the effects increase monotonically when using within-industry emission intensity quartiles. We also develop an alternative, rule-based keyword approach to validate and fine-tune the green-deal classification of ClimateBERT and show that the results hold. We investigate to what extent our results from the analyses requiring non-missing financial data from Orbis could suffer from selection bias and do not find evidence that this is the case. Finally, we investigate heterogeneity of the findings along the sectoral and time dimension and show that the results are consistent with a causal effect from regulatory tightening in the EU ETS driving surges in the carbon price.

This paper contributes to at least three interconnected strands of the corporate finance literature: the effects of climate change regulation on corporate behaviour, the strategic use of mergers and acquisitions, and the role of financial frictions in environmental adaptation.

First, we add to the rapidly growing literature on corporate finance and climate risk. While a substantial body of work explores how transition risks and carbon emissions are priced in financial markets or influence institutional investors (Bolton and Kacperczyk 2021; Ilhan, Sautner, and Vilkov 2021; Pastor, Stambaugh, and Taylor 2022), there is comparatively less evidence on the operational adjustments firms make when faced with sudden, binding regulatory shocks. Existing studies highlight that firms may react to environmental pressures through internal abatement or shifting emissions across borders (e.g., Bartram, Hou, and Kim 2022; Colmer et al. 2025). We contribute to this literature by exploiting an

unexpected tightening of the EU ETS to show that firms actively improve their emission efficiency, by emphasizing the role of mergers and acquisitions for firms' decarbonization, and by documenting significant heterogeneity across sectors: power producers achieve deeper internal emission reductions than manufacturing firms, consistent with the latter facing steeper marginal abatement costs.

Second, we contribute to the M&A literature by identifying a novel, climate-driven motive for corporate takeovers. Traditional M&A literature establishes that firms frequently use acquisitions to acquire new technologies, externalize R&D, or rapidly adapt to industry-wide shocks when internal innovation is too slow or costly (e.g., Bena and Li 2014; Harford 2005; Mitchell and Mulherin 1996; Phillips and Zhdanov 2013). We lay out a similar mechanism for firms' reaction to climate transition risk. We show that manufacturing firms that are most affected by the regulatory shock are more likely to change their acquisition strategy and focus on green firms. We also show that green takeovers are associated with a significant reduction in emissions.

Finally, our paper is also related to recent work on financial frictions and decarbonisation. A growing literature studies whether and how financial constraints impede decarbonization (Bartram, Hou, and Kim 2022; Boeckx, Struyfs, and Torsin 2025; De Haas et al. 2025) and finds mixed results that are context- and sample-specific. We find no evidence that financial constraints drive the observed responses in green M&A activity and attribute this to the benign macroeconomic environment characterized by highly accommodative monetary policy and ample market liquidity. Under these conditions, external financing costs were likely sufficiently low that even relatively constrained firms could secure external finance to fund their acquisition.

2. INSTITUTIONAL BACKGROUND

Established in 2005, the EU ETS regulates greenhouse gas emissions of power producers, energy-intensive industries, aviation (since 2013), and maritime transport (since 2024). A stationary installation is covered by the EU ETS if it performs a regulated activity and surpasses an activity-specific capacity threshold as set out in EU Directive 2003/87/EC. The EU ETS sets an annual cap on total emissions and requires regulated installations to surrender emission allowances (EUAs) for each tonne of CO₂-equivalent emitted in a given year. Participants in the EU ETS can obtain EUAs through auctions, trading in the secondary market, or, for some industrial sectors, free allocation by the regulator.

The EU ETS is structured in discrete trading phases designed to allow for frequent regulatory updates (Sato et al. 2022). Phase I (2005–2007) was designed as a pilot phase, followed by Phase II (2008–2012) which corresponded to the First Commitment Period of the Kyoto Protocol. Our analysis focuses on Phase III (2013–2020) and the first years of Phase IV (2021–2030). Phase III was characterized by the harmonization of the system’s administration which up to this point had consisted of linked but largely independent systems administered by the member states (Ellerman, Marcantonini, and Zaklan 2016). Although Phase III addressed important concerns, e.g., the uncertainty of the annual allowance cap (Ellerman, Marcantonini, and Zaklan 2016) or windfall profits due to grandfathering of allowances based on historical emissions (Sato et al. 2022), the system was held back by a structural oversupply of allowances that depressed carbon prices and cast doubt on the system’s ability to set a credible long-term price signal (Sitarz et al. 2024).

To address the oversupply of allowances, the EU Commission backloaded auctioning of allowances from 2014–2016 to later years and even proposed a permanent mechanism, the so-called market stability reserve (MSR) to withhold allowances and introduce short-term

scarcity in the system. Because the MSR would only reduce allowance supply in the short term without affecting their long-term availability over the system's lifetime, carbon prices did not respond meaningfully, consistent with intertemporal arbitrage (Perino and Willner 2016).

Figure 1 shows that the price dynamics in the EU ETS changed fundamentally in 2017 after the unexpected inclusion of an invalidation rule which would automatically cancel excess allowances in the MSR³ (Borghesi et al. 2023) thereby credibly reducing the long-term supply of allowances and causing allowance prices to quadruple from around EUR 6 in 2016 to EUR 25 before the end of 2018 (Bruninx, Ovaere, and Delarue 2020).⁴ While the MSR reform succeeded in restoring faith in the EU ETS, carbon prices quadrupled once more in the context of the announcements of the Climate Law in 2020 and the Fit-for-55 package in 2021 (see Figure 2).

Figures 1 and 2 HERE

Regulated emissions in the EU ETS have declined significantly over the sample period. Verified emissions totaled ± 1.75 billion tonne CO₂ equivalent in 2013 and have gone down to ± 1.1 billion by 2023. As can be seen in Figure 3, this decline picked up speed after 2017 and is more pronounced in the power sector than in manufacturing.

Figure 3 HERE

3. DATA

We combine data from various sources to construct a firm-year panel. First, the EU ETS Union Registry ⁵ records data of installations, e.g., power plants or steel mills, regulated by

the EU ETS. The EU ETS Union Registry reports installation-level data on annual verified emissions, emission allowances allocated for free, industry, main activity, location of the installation, and owner. Regulated installations are linked to a unique Operator Holding Account which in turn can be linked to a firm via the matching provided by Cameron and Ho (2024). In this way, we can aggregate emissions from the installation-year level to the firm-year level, which we then match to annual balance sheet and P&L data from Bureau van Dijk's Orbis dataset that we clean following the methodology proposed by Kalemli-Özcan et al. (2024).

Our final sample consists of 24,202 firm-year observations from 2,207 firms that own 5,314 regulated installations over the period 2013–2023. Table 1 presents summary statistics of the main variables for the full sample and separately for power and manufacturing. On average, firms own 2 installations emitting more than 500,000 tCO₂ and generate more than 900 million EUR operating revenue annually. Emissions and operating revenue are highly skewed with the largest firms having five (two) times higher annual emissions (operating revenue) than the average. Firms' size in terms of number of installations, emissions, and operating revenue is also unevenly distributed with power producers being materially larger than manufacturing firms across all dimensions. Figure 3 shows that our sample covers approximately 90% of total verified emissions in the EU ETS.⁶ Firm-year observations with valid Orbis balance sheet data cover between 65 to 70% of total verified emissions in the EU ETS.⁷

In a second step, we collect data on M&A deals announced by firms in our sample over the period 2013–2023 from Bureau van Dijk's Zephyr database and keep those that were completed at the latest by January 2025 (the date at which we collected this data). In addition to identifier information on the acquirer and target, Zephyr also provides textual information about the rationale for the deal as given by the acquiring firm. We leverage the pretrained

large language model ClimateBERT (Webersinke et al. 2022) on the textual information about the deal rationale to classify an M&A deal as ‘green’ or not.⁸ We hand-check the validity of this approach and give some examples in Appendix A. In the robustness section we also cross-validate the classification of ClimateBERT with an alternative keyword-based classification and analyze the robustness of our findings using this alternative classification.

Our sample covers 2,834 deals completed by 601 firms between 2013–2023 averaging around 260 deals per year of which approximately 20% (547 out of 2,834) are classified as green by ClimateBERT. Figure 8 shows that while the annual number of deals remains stable over the sample period, the annual number of green deals triples from less than 30 at the start of the sample period to more than 90 in 2023.

Table 1 HERE

4. EMPIRICAL SETUP AND IDENTIFICATION STRATEGY

4.1. *Difference-in-Differences Setup*

We exploit the regulatory tightening of the EU ETS as a natural experiment to compare firms with varying exposure to the steep carbon price increase. Following Mulier, Ovaere, and Stimpfle (2024), we construct a time-invariant measure of a firm’s emission intensity which we then use to proxy for firms’ carbon price exposure. We construct this as the emission-weighted average emission intensity of the installations owned by a firm over the period 2013–2016. Specifically, for firm f owning installations $j_1, \dots, j_n \in I_f$, we set

$$\text{Intensity}_f = \sum_{j \in I_f} \omega_j \times \text{Intensity}_j, \quad (1)$$

where

$$\omega_j = \frac{\overline{\text{Emissions}}_j}{\sum_{j \in I_f} \overline{\text{Emissions}}_j}, \quad (2)$$

and where $\overline{\text{Emissions}}_j$ are the average verified emissions of installation j in years 2013–2016 and Intensity_j is the emission intensity of installation j .

We infer an installation’s emission intensity in one of two ways depending on the type of installation. First, for power producers we match power plants with the Global Energy Monitor database⁹ and approximate the plant’s emission intensity in tCO₂ per unit of generated electricity by its fuel type (e.g., bituminous coal, lignite, natural gas) and technology (e.g., steam turbine, combined cycle) (Nicholson and Heath 2021). Second, for all other installations covered by the EU ETS Free Allocation Methodology¹⁰ we approximate an installation’s production-based emission intensity as follows: An installation’s free allocation of EU allowances is proportional to its so-called historical activity level measured in units of product (e.g., tonnes steel) multiplied by a product-specific efficiency benchmark. Because we observe an installation’s main activity and its free allocation in the EU ETS Union Registry, we can back out its historical activity level to get a measure of actual production. Our measure of production-based emission intensity is the ratio of observed historical emissions and inferred historical activity level. Note that this measure is time-invariant and based on the median emissions and production over the so-called baseline period specified in the Free Allocation Methodology (which equals 2005–2008 or, where higher, 2009–2010).¹¹

The higher a firm’s emission Intensity_f , the more emission allowances it has to surrender per unit of production and the more exposed it is to a marginal increase in carbon prices relative to more efficient firms. Our goal is to compare firms producing the same products (e.g., electricity, steel, or cement) that mainly differ in the cleanliness of their production.

To this end, we define a firm f as high-exposed if its emission Intensity $_f$ is above the median emission intensity in its NACE 4-digit industry i , i.e.,

$$\text{Exp}_f = \begin{cases} 1, & \text{if Intensity}_f > \text{median}_i^{\text{intensity}}, \\ 0, & \text{else,} \end{cases} \quad (3)$$

To estimate the effect of the tightening in climate regulation on relatively more exposed firms, we use the following Difference-in-Differences specification:

$$Y_{f,t} = \beta_1 \text{Post}_t + \beta_2 \text{Exp}_f + \beta_3 \text{Post}_t \times \text{Exp}_f + X_{f,t} + \epsilon_{f,t}, \quad (4)$$

where $Y_{f,t}$ is the outcome variable of interest (e.g., emissions or emissions over operating revenue), Post_t is a dummy equal to one if $t \geq 2017$ and zero otherwise, and $X_{f,t}$ are a set of control variables. Our main coefficient of interest is β_3 which captures the effect of the regulatory tightening on high exposed relative to low exposed firms. We cluster standard errors at the firm level across all specifications. We also estimate versions of equation (4) with year fixed effects γ_t and firm fixed effects α_f or, alternatively, year-country $\nu_{t,c}$ and year-industry $\mu_{t,i}$ fixed effects. When including fixed effects the parameters β_1 and β_2 will be omitted.

To assess the dynamic effects of the regulatory tightening we furthermore estimate event study versions of equation (4)

$$Y_{f,t} = \sum_{s=2013}^{2023} \beta_s \mathbb{1}_t^s \times \text{Exp}_f + X_{f,t} + \alpha_f + \nu_{t,c} + \mu_{t,i} + \epsilon_{f,t}, \quad (5)$$

where $\mathbb{1}_t^s$ is a dummy that equals one if year t equals s , and 0 otherwise. We generally use 2016 as the reference year, i.e., omit it from equation (5) so that the estimated coefficients $\hat{\beta}_s$

show the difference between high-exposed compared to low-exposed firms relative to 2016, the last year before the regulatory tightening.

Following Dechezleprêtre, Nachtigall, and Venmans (2023) and Mulier, Ovaere, and Stimpfle (2024), we estimate the proportional average treatment effect (Chen and Roth 2024) by applying the Poisson Quasi-Maximum Likelihood (QML) estimator (Wooldridge 2023) to equation (4) which does not require any distributional assumptions beyond the specification of the conditional mean and does not suffer from the incidental parameter problem (Wooldridge 2023). Because our main outcome variables of interest (emissions and emissions over operating revenue) are non-negative, the exponential conditional mean implied by Poisson QML is a natural choice with two appealing features: First, coefficient estimates can be interpreted as percentage changes in the outcome variable. Second, Poisson QML implies a parallel trends assumption in the ratio of means which is more credible than parallel trends in levels in our context of firms with heterogeneous emission intensities. In our analysis of firms' propensity to engage in (green) deals, we estimate equation (4) via OLS as our baseline estimate.¹²

4.2. Identifying Assumptions

Our main identifying assumption is that in the absence of the regulatory tightening, relatively high-exposed and relatively low-exposed firms would have followed parallel trends. While we cannot empirically test this assumption, its plausibility is supported by two observations.

First, structural oversupply of emission allowances (Sato et al. 2022) combined with lacking long-term credibility of the EU ETS (Sitarz et al. 2024) caused carbon prices to remain at depressed levels of around 5 to 10 EUR/tCO₂, which were generally considered to be too low to meaningfully affect firms' abatement behavior (Fuss et al. 2018). Although the EU

Commission proposed the so-called market stability reserve (MSR) as a mechanism to manage the system's oversupply already in 2015, carbon prices only reacted substantially after the unexpected inclusion of the cancellation mechanism, which for the first time credibly reduced the long-term supply of emission allowances and caused allowance prices to quadruple within less than two years (Bruninx, Ovaere, and Delarue 2020).

Second, a firm's exposure to the regulatory tightening is defined within a narrow NACE 4-digit industry which enables us to compare firms producing similar products, operating in the same markets, facing similar product demand, which are therefore unlikely to deviate from parallel trends in the absence of the regulatory tightening. We can empirically assess the plausibility of parallel trends before the regulatory tightening by estimating event study versions of our main specification (Equation (5)).

4.3. (Local) Average Treatment Effects

While firms' responses to the EU ETS have been studied before, they often exploit capacity-specific inclusion criteria to compare firms falling under the EU ETS to similar, unregulated firms. As a consequence, existing work typically estimates a local average treatment effect for firms close to the threshold. For example, in a matched difference-in-differences design Dechezleprêtre, Nachtigall, and Venmans (2023) and Colmer et al. (2025) both find that the introduction of the EU ETS in 2005 caused firms above the inclusion threshold to reduce emissions relative to firms below. Generalizations of these local treatment effect estimates to the larger regulated firms which are further above of the threshold or to more recent periods are difficult because in many industries the EU ETS covers virtually all emitters.¹³ In contrast, our empirical strategy exploits variation in the efficiency of firms within the EU ETS which enables us to study the vast majority of firms regulated by the EU ETS and the most recent regulatory changes. It also means that our study is not limited to a set of countries for which emission data for non-regulated entities is available (France in Colmer

et al. (2025), France, Netherlands, and Norway in Dechezleprêtre, Nachtigall, and Venmans (2023)).

5. EMPIRICAL ANALYSIS

We present four sets of results. First, we document that the consecutive regulatory tightening in the EU ETS reduced the emissions of the high-exposed firms significantly relative to the low-exposed firms. Second, we show that this led to improved emission efficiency as the reduction in emissions is not driven by a reduction in operating revenue. Third, we investigate the role of M&A activity in firms' abatement behavior. Our results show that the high-exposed firms became more likely to acquire green firms after the regulatory tightening relative to low-exposed firms, and show suggestive evidence that this mattered for the observed reduction in emissions. Finally, we present a number of tests that ensure the robustness of our main findings.

5.1. *Regulatory Tightening and Emissions*

We first document the impact of the regulatory tightening on firms' emissions. We start with insights from descriptive statistics and figures, before turning to regression analyses. We graphically compare the average emissions, expressed in tCO₂, of high-exposed and low-exposed firms during the sample period. We do this separately for firms in the power sector and manufacturing industries (Figures OA.2 and OA.3 in the Online Appendix).

For firms in either the power sector or the manufacturing industry, it is the case that high-exposed firms emit significantly more than low-exposed firms. This is partly by design but also related to firm size (which we will absorb in the regression with firm fixed effects). The average high-exposed power firm emitted around 4 million tCO₂ before the regulatory tightening, roughly 5 times more than the average low-exposed power firm. The average

high-exposed manufacturing firm emitted around 450 thousand tCO₂ before the regulatory tightening, roughly 3 times more than an average low-exposed manufacturing firm. A difference in the level of emissions does not violate our identification strategy. What does matter is that the emissions of both high- and low-exposed firms are evolving similarly before the regulatory tightening (i.e. have parallel trends), which is exactly what the right panel of these figures suggest.

The descriptive statistics further show that it is particularly the high-exposed power firms that have reduced their emissions over the sample period, and that power firms in general were already reducing their emissions before the regulatory tightening. To understand the latter, it is important to note that abatement decisions in the power sector are equally driven by relative fuel prices and carbon prices (Cullen and Mansur 2017). In the regression analyses, we will therefore add control variables in order to isolate the effect of the carbon price (i.e. the regulatory tightening) from that of relative fuel prices. An event study analysis will confirm that, holding relative fuel prices constant, there are no discernible pre-trends in the power sector. Regarding the manufacturing firms, the descriptive statistics very clearly show material deviations after the tightening where the high-exposed firms kept emissions stable for two years and then started reducing and low-exposed firms even slightly increased emissions for two to three years and then started reducing.

Next, we move to the regression analyses. Table 2 presents Poisson QML estimates of equation (4) with emissions as the dependent variable. Column 1 presents estimates of equation (4) without fixed effects. The estimate for Post_t is negative but statistically insignificant, indicating that emissions of low-exposed firms decreased after the regulatory tightening but not significantly. The estimate for Exp_f indicates that high-exposed firms on average have 270% ($e^{1.306} - 1$) higher (or 3.7 times more) emissions than low-exposed firms before the regulatory tightening, in line with the descriptive statistics.

Table 2 HERE

Our main coefficient of interest is β_3 corresponding to $\text{Post}_t \times \text{Exp}_f$ which captures the impact of the regulatory tightening on high-exposed relative to low-exposed firms. We find that the consecutive regulatory tightening events are associated with an average emission reduction of around 19.3% ($e^{-0.214} - 1$). Column 2 indicates that this estimate is not meaningfully affected by the inclusion of year and firm fixed effects. Because power producers' marginal production decisions depend on relative fuel prices (Cullen and Mansur 2017) and the availability of zero marginal cost sources (e.g., solar or wind) (Fell and Kaffine 2018), we include the difference between the dark spread and spark spread, i.e., the difference in theoretical gross margins of selling a unit of electricity produced from coal (dark spread) or gas (spark spread), and the country-level annual average share of renewable energy as control variables, both interacted with the exposure dummy Exp_f . We moreover control for the UK's Carbon Price Support, a carbon tax for UK power producers introduced in 2013 (Gugler, Haxhimusa, and Liebensteiner 2023; Leroutier 2022; Abrell, Kosch, and Rausch 2022). While the estimated effect reduces, it remains economically and statistically significant at around 10.4% ($e^{-0.11} - 1$). In column 4, we flexibly account for other time-varying heterogeneity at the country or industry level by additionally including year-country and year-industry fixed effects and find emissions reductions of 13% ($e^{-0.139} - 1$).

Because fuel switching is an abatement channel unique to power producers, we perform a sample split to investigate the heterogeneity of emission abatement between power and manufacturing. Column 5 shows that power producers achieved much higher emission reductions of around 32% ($e^{-0.391} - 1$) compared to more modest, but significant reductions of around 5% ($e^{-0.052} - 1$) in manufacturing. This heterogeneity is likely driven by power producers' ability to reduce carbon costs by switching production from emission-intensive coal-fired to cleaner gas-fired power plants (Cullen and Mansur 2017) and is in line with estimates of

lower marginal abatement costs for power producers (Duscha et al. 2022). In the robustness section, we dig deeper in possible heterogeneities within manufacturing subsectors.

A causal interpretation of our estimates in columns 4, 5, and 6 of Table 2 crucially hinges on the parallel trends assumption. The descriptive statistics suggest that this assumption likely holds, but to assess its plausibility empirically, we estimate equation (5), which is an event study version of equation (4).

The event study coefficients β_s are estimated relative to 2016, the last year before the announcement of the regulatory tightening, and enable us to empirically assess the annual difference in emissions between high-exposed and low-exposed firms. Rather than showing the regression output in table format, we show the estimated $\hat{\beta}_s$ coefficients and their 90% confidence intervals in Figures 4 and 5. Figure 4 shows the results for the full sample (i.e. the event-study equivalent of column 4 of Table 2). Figure 5 shows the results for the power sector in panel (a) and manufacturing industry in panel (b) (i.e. the event-study equivalent of, respectively, column 5 and 6 of Table 2)

Figures 4 and 5 HERE

Three interesting conclusions can be drawn from these figures. First, before the regulatory tightening, the coefficients are small and statistically not different from zero, indicating that high- and low-exposed firms were on parallel trends, and this in both the power and manufacturing sample. Second, after the announcement of the first regulatory tightening in 2017, high-exposed firms start to significantly reduce emissions relative to low-exposed firms in the same NACE 4-digit industry in both the power and manufacturing sample. Third, the effect is larger for power producers than for manufacturing firms (Figure 5). Additionally, in both sectors the effect becomes economically stronger over time with a strengthening of the

effect after the announcement of the second regulatory tightening in 2020, which is in line with the increase in the carbon price that is associated with these regulatory tightenings (see Figure 2). By 2023, high-exposed power producers have reduced their emissions by 65.4% ($e^{-1.061} - 1$) relative to low-exposed power producers. For high-exposed manufacturing firms, there appears to be an immediate effect on emissions of around 4-5% that persists until 2020; followed by an additional reduction of around 2-3%, increasing the relative reduction in emissions to up to 7.5% ($e^{-0.078} - 1$) as of 2021.

5.2. ETS Regulation: Impact on Firms' Emission Efficiency

The previous section focused purely on emission abatement of firms as a response to the regulatory tightening. These emission reductions could reflect either genuine efficiency improvements or mere output contraction. To distinguish between these mechanisms, we examine whether high-exposed firms reduce economic activity, are able to maintain, or even manage to expand their economic activity relative to low-exposed firms after the regulatory tightening, where our prime measure of activity is firms' operating revenue. We also directly investigate whether high-exposed firms' emission efficiency changes after the regulatory tightening relative to low-exposed firms, where emission (in)efficiency is the ratio of emissions (in tonne CO₂ equivalent) to operating revenue (in thousand euros).¹⁴

Table 3 presents difference-in-differences estimates like equation (4) having either operating revenue or emission efficiency as dependent variable. Columns (1)-(3) show that the regulatory tightening had no significant effect on firms' operating revenue. The coefficient on $\text{Post}_t \times \text{Exp}_f$ is close to zero and statistically insignificant in all three samples (full sample, power producers, and manufacturing firms), indicating that high-exposed firms do not reduce output relative to low-exposed firms following the regulatory shock. This null result on operating revenue provides initial evidence against the output contraction hypothesis.¹⁵

Table 3 HERE

Turning to estimates of emission efficiency gains in columns (4)-(6) we find that high-exposed firms achieve an 8.5% improvement in emission efficiency after the regulatory tightening (significant at the 1% level). This aggregate effect masks substantial sectoral heterogeneity: power producers improve efficiency by 21.2% ($e^{-0.239} - 1$), while manufacturing firms show more modest but still economically meaningful gains of 5.4% ($e^{-0.056} - 1$). These magnitudes are consistent with our absolute emission results, and indicate that the emission reductions are not merely coming from proportional output adjustments but seem to translate into genuine efficiency improvements.

Following the same structure as in the previous section, we also zoom in on temporal dynamics in emission efficiency using event-study specification equation (5). Figure 6 presents results for the full sample. The pre-period coefficients satisfy the parallel trends assumption, followed by efficiency improvements that strengthen through 2020 (to around 12%) before moderating in the final years of our sample, with the 2023 coefficient indicating a 8% ($e^{-0.084} - 1$) efficiency improvement for high-exposed firms relative to low-exposed firms. Note, also, that while the confidence intervals widen in 2021 to 2023, we still find statistically significant efficiency improvements (except for 2022). Moreover, the estimated coefficient for 2021, 2022, and 2023 are still larger than the immediate effects of 2017 and 2018.

Figure 6 HERE

Sectoral heterogeneity in efficiency dynamics is pronounced, as shown in Figure 7. Power producers (panel a) exhibit the strongest response, with effects intensifying dramatically after 2017 to reach approximately -0.56 by 2023, which implies a 43% improvement in emission efficiency. The results for manufacturing firms (panel b) are displaying parallel trends in the

pre-period and gradual efficiency improvements emerging after 2017 that strengthen over time until 2020 and then weaken and become statistically insignificant as of 2021 (due to a reversal of the effect and much larger standard errors).

Figure 7 HERE

Summing up, the null effects on operating revenue and the significant effects on emission efficiency provide evidence that firms achieved lower emissions through operational adjustments rather than output contractions. Indeed, if high-exposed firms were simply downsizing in response to higher carbon costs, we would observe coordinated reductions across emissions, operating revenue, employment, and assets. Instead, we find that only emissions consistently decline, while real activity measures remain stable. Only for high-exposed manufacturing firms, it appears that the picture is not 100% clear. While we do observe a significant improvement in emission efficiency over the post-period on average, and while we do not observe a significant reduction in operating revenue over the post-period on average, we cannot rule out no significant difference in emission efficiency by the end of our sample period.

5.3. The Market for Corporate Control

In this section, we focus on the market for corporate control as a potential channel through which firms restructure their operations in response to higher carbon costs. M&As represent a particularly relevant margin of adjustment: they allow firms to rapidly acquire new technologies, enter cleaner production processes, or access complementary assets that reduce emissions (Mitchell and Mulherin 1996; Harford 2005). Moreover, compared to investments in R&D or gradual technology adoption, acquisitions offer a faster path to strategic repositioning when firms face acute regulatory pressure.

We examine five specific questions. First, do high-exposed firms increase their overall acquisition activity after the regulatory tightening relative to low-exposed firms? Second, do high-exposed firms shift toward acquiring green targets, i.e. firms with lower carbon footprints or cleaner technologies? Third, are these green acquisitions associated with actual emissions reductions? Fourth, do financial constraints limit firms' ability to pursue this strategic response? And finally, do firms engage in regulatory arbitrage by increasing acquisitions of targets located outside the EU ETS jurisdiction, which would suggest carbon leakage through the market for corporate control.

5.3.1. Propensity to Engage in (Green) Acquisitions

We begin by examining whether high-exposed firms increase their overall M&A activity following the regulatory tightening, before turning to the composition of these deals. Throughout this analysis, we use announcement dates rather than completion dates to capture the timing of strategic decisions, though we restrict our sample to deals that were ultimately completed by January 2025. This ensures we measure actual acquisition decisions while avoiding the mechanical lag between announcement and completion.

The left panel of Figure 8 presents the number of M&A deals announced by firms in our sample over the period 2013–2023, separately for high-exposed and low-exposed firms. On average, high-exposed firms announce more deals than low-exposed firms, but the number of M&A deals evolves similarly between the two groups over time, both before and after the regulatory tightening.

Figure 8 HERE

To assess firms' propensity to engage in M&A deals more rigorously, we estimate equation (4) where the dependent variable $Y_{f,t}$ is the number of acquisitions that firm f announces in year

t . Column 1 of Table 4 shows that the coefficient on Post_t is close to zero and statistically insignificant, indicating that low-exposed firms do not announce more M&A deals after the regulatory tightening. The coefficient on Exp_f is positive and statistically significant, indicating that high-exposed firms announce more M&A deals than low-exposed firms before the regulatory tightening. The coefficient of interest on $\text{Post}_t \times \text{Exp}_f$, however, is close to zero and statistically insignificant, indicating that high-exposed firms do not start to announce more M&A deals after the regulatory tightening relative to low-exposed firms. In other words, while high-exposed firms do more deals than low-exposed, this difference remains stable after the regulatory tightening. Columns 2 to 6 of Table 4 show that this result is robust to controlling for increasingly tighter fixed effects and holds for both the power and manufacturing samples.

Table 4 HERE

We next analyze firms' propensity to engage in green M&A deals using the same diff-in-diff setting as before. As explained in Section 3, we leverage the pretrained large language model (LLM) ClimateBERT to label a deal as "green" based on the deal rationale given by the acquiring firm.¹⁶ The right panel of Figure 8 provides descriptive evidence of an increase in the number of green deals over the entire sample period, and substantially stronger after the regulatory tightening. The figure also shows that before the regulatory tightening, high-exposed and low-exposed firms were equally likely to engage in green deals, but that after the regulatory tightening, high-exposed firms announced more green deals than low-exposed firms.

We test the effect of the regulatory tightening on firms' propensity to engage in green deals by estimating equation (4) with the number of green deals that firm f announces in year t as the dependent variable. Column 1 of Table 5 confirms the descriptive evidence presented

in Figure 8: First, before the regulatory tightening, high-exposed and low-exposed firms were equally likely to engage in green deals. Second, firms' propensity to engage in green deals increases significantly after the regulatory tightening. Third, the increase in green deal propensity is significantly stronger for high-exposed firms. This finding is robust to additional control variables (column 3) and tighter fixed effects (column 4). Columns 5 and 6 show that the increase in green deals is driven by firms in the manufacturing sector, with no significant change in the power sector.

Table 5 HERE

Figure 9 presents event study estimates separately for power producers (Panel a) and manufacturing firms (Panel b). We statistically confirm the descriptive evidence on the parallel trends assumption in Figure 8 using equation (5) on the number of 'green' M&A deals. In both sectors, the estimated coefficients are close to zero and statistically insignificant in the period before the regulatory tightening, confirming that the propensity to announce green deals for high- and low-exposed firms evolved similarly before 2017. The post-tightening dynamics, however, differ markedly across sectors. For power producers, the coefficients remain statistically insignificant throughout the post-period, consistent with the null finding from the difference-in-differences specification in Table 5. In contrast, for manufacturing firms, we observe a gradual and persistent increase in green M&A activity among high-exposed firms relative to low-exposed firms. This pattern indicates that the significant effect documented in the difference-in-differences framework emerges from a progressive strengthening of the green acquisition response over time rather than an immediate discrete jump. The gradual nature of this response likely reflects the time required for manufacturing firms to identify suitable green targets and complete strategic acquisitions in response to sustained regulatory pressure.

Figure 9 HERE

5.3.2. *Green Acquisitions and Emission Reductions*

Having established that high-exposed firms increase their propensity to engage in green M&A activity following the regulatory tightening, we next examine whether these green acquisitions are associated with actual emissions reductions. This analysis addresses a distinct but complementary question: while our main results demonstrate the effect of regulation on green M&A activity, it is important to establish that these green acquisitions are associated with meaningful emission reductions. To this end, we analyze the emissions trajectory of firms before and after they complete green acquisitions.

Table 6 presents results from a within-firm analysis where we regress emissions on an indicator for having completed at least one green deal by year t . This specification compares the same firm's emissions before and after engaging in green M&A, controlling for firm fixed effects and progressively more demanding time-varying effects. Columns (1)-(2) presents results without firm controls, while columns (3)-(4) include log operating revenue and log total assets to account for scale effects (and thus results again in a smaller sample). All specifications include firm fixed effects, with columns (2)-(4) adding year-country and year-industry fixed effects to absorb common time-varying shocks across countries and sectors. Column 4 further includes a post-period interaction to test whether the emissions-reduction effect strengthened after the 2017 regulatory tightening.

Table 6 HERE

The results provide consistent evidence that green M&As are associated with lower emissions. The negative association is robust across all specifications with progressively demanding fixed effects, and the magnitude is economically meaningful—consistent with efficiency gains

rather than mere scale adjustments, as firm size controls do not eliminate the emissions reduction. In our most demanding specification without the post interaction (column 3), we find an 11.1% ($e^{-0.118}-1$) reduction in emissions associated with green deals. We do not find that effect differs significantly post-tightening (column 4).

As a complementary test of the link between green M&As and emissions outcomes, we assess the counterfactual emission reductions that would have occurred absent green acquisitions. Specifically, we re-estimate our main emission regressions (columns 4-6 of Table 2) after excluding all firms that completed at least one green deal during the sample period. The results, reported in Online Appendix Table OA.2, show that excluding firms engaged in green M&A weakens the observed emission reduction by approximately 5 percentage points for high-exposed power producers and 1 percentage point for high-exposed manufacturing firms—representing roughly 16% of the total emission reduction in both sectors.¹⁷

This evidence further supports the role of the market for corporate control as a meaningful channel through which firms adjust to regulatory tightening. Combined with our main findings, these results actually provide evidence of the complete transmission mechanism: regulatory tightening induces increased green M&A among exposed firms, which in turn is associated with meaningful emissions reductions.

5.3.3. Financial Constraints

After documenting that high-exposed firms increase their green M&A activity and that these acquisitions are associated with emissions reductions, we next investigate whether financial constraints moderate this strategic response.

Financial constraints might limit firms' ability to pursue acquisitions in response to regulatory shocks. To examine this, we construct three ex-ante proxies for firms' financing

capacity, measured over the pre-treatment period from 2013 to 2016: cash holdings (ratio of cash and cash equivalents to total assets), leverage (ratio of financial debt to total assets), and firm size (log of total assets). For each proxy, we create a dummy variable indicating firms as being financially constrained if they are in the bottom tercile for cash and size, or top tercile for leverage. We then estimate triple interaction specifications to test whether financially-constrained, high-exposed firms respond differently to the regulatory tightening in terms of green M&A activity than financially-unconstrained, high-exposed firms.

Table 7 presents results from these triple interaction specifications. Across all three financial constraint measures, we find that the triple interaction coefficient is small in magnitude and statistically insignificant. These null results suggest that traditional measures of financial constraints do not materially moderate the green M&A response of high-exposed firms to the regulatory tightening.

Table 7 HERE

We interpret these insignificant findings in the macro-economic context during our sample period. The period 2013-2023 is characterized by exceptionally accommodative monetary policy by the European Central Bank, with negative policy rates and ample liquidity provision. This environment meant that even financially constrained firms had relatively good access to external capital, potentially masking the role of financial constraints that might emerge in tighter financing environments. Furthermore, if green M&As are strategically important and are relatively modest in scale, then they could also be financed from operating cash flows.

5.3.4. *Acquisitions and Carbon Leakage*

An important concern in climate policy is that firms may respond to domestic carbon pricing by relocating production to jurisdictions with weaker environmental regulation, also known as carbon leakage. In the context of M&A activity, this would imply that high-exposed firms increasingly acquire targets located outside the EU ETS area, allowing them to shift production away from regulated facilities while maintaining overall output. Such regulatory arbitrage would undermine the environmental effectiveness of the policy and suggest that observed emissions reductions within the EU ETS reflect displacement rather than genuine abatement.

To test for this channel, we estimate equation (4) using as the dependent variable the number of deals in which the target firm is located outside EU ETS member states. Table 8 shows that high-exposed firms do not increase their acquisition of non-EU ETS targets following the regulatory tightening. This result suggests that there is no carbon leakage via the market for corporate control.¹⁸

Table 8 HERE

5.4. *Robustness*

Our baseline results establish that the regulatory tightening in the EU ETS induced significant emission reductions among high-exposed firms and shifted their acquisition strategies toward green targets. To ensure that our findings are not artifacts of particular measurement or estimation decisions, we run three distinct classes of robustness tests. First, we examine whether our binary exposure classification masks important nonlinearities or discards meaningful variation in treatment intensity. Second, we probe the stability of our emission effects across manufacturing subsectors and investigate whether sample composition issues

related to Orbis coverage or contemporaneous shocks could drive our results. Third, we validate our key methodological choices regarding the estimator and the classification of green acquisitions.

5.4.1. Treatment intensity: continuous exposure measure

Our baseline specification uses a binary exposure indicator (Exp_f) that equals one for firms with above-median emission intensity within their 4-digit NACE industry. While this approach offers clear interpretation -comparing high-exposed to low-exposed firms before and after the regulatory tightening- it imposes a discrete treatment threshold that may mask heterogeneity and discard information about the intensity distribution. We therefore examine two alternative exposure definitions: a continuous intensity measure and a quartile-based classification.

Table OA.3 in the Online Appendix presents results using the continuous emission intensity measure ($Intensity_f$) instead of the binary exposure dummy. To facilitate interpretation, we standardize $Intensity_f$ to have zero mean and unit standard deviation. The coefficient on $Post_t \times Intensity_f$ captures how emission reductions scale with pre-treatment intensity. Results confirm our main findings: a one-standard deviation higher emission intensity is associated with an 11.1% ($e^{-0.118} - 1$) greater emission (per standard deviation change) reduction after 2017. Consistent with our baseline estimates, the continuous-treatment specification shows stronger effects in the power sector (18.0% reduction per standard deviation) than in manufacturing (4.8% reduction per standard deviation).

The continuous approach reveals that our findings are not an artifact of the binary classification; emission reductions scale monotonically with pre-treatment intensity. Moreover, the magnitudes are economically comparable to our baseline estimates, suggesting the binary split captures the essential variation in treatment intensity without substantially distorting

effect sizes.

We further examine whether effects concentrate in the upper tail of the intensity distribution by estimating quartile-based specifications (Table OA.4 in the Online Appendix). Defining quartile dummies within 4-digit NACE industries and omitting the first quartile as the reference group, we find that emission reductions are not confined to the highest-intensity firms. Relative to first-quartile firms, those in the second, third, and fourth quartiles all exhibit statistically significant emission reductions. Effects increase mostly monotonically across quartiles. This pattern confirms meaningful treatment effects across much of the intensity distribution, not just at the extremes.

Taken together, these alternative specifications demonstrate that our core finding—greater emission reductions by higher-intensity firms following regulatory tightening—is robust to how we measure exposure and does not depend on the binary median split used in our baseline analysis.

5.4.2. Manufacturing subsector heterogeneity

Our baseline results reveal substantially larger emission reductions among power producers than manufacturing firms, consistent with lower marginal abatement costs and the availability of fuel switching options in electricity generation. However, manufacturing encompasses diverse production technologies with potentially heterogeneous abatement opportunities, ranging from energy-intensive basic materials production (steel, cement, chemicals) to lighter manufacturing activities. Understanding where emission adjustments concentrate within this heterogeneous sector is important both for interpreting our aggregate manufacturing effects and for fine-tuning climate policies.

Table OA.5 in the Online Appendix disaggregates our manufacturing sample by NACE

2-digit industry codes, isolating three major emission-intensive subsectors: basic metals (NACE 24, primarily steel), non-metallic minerals (NACE 23, primarily cement), and chemicals (NACE 20). All remaining manufacturing industries are grouped in an “all other” category. The results reveal meaningful heterogeneity in emission responses. Basic metals show the strongest response with a coefficient of -0.096 (significant at the 1% level), nearly twice the magnitude of our aggregate manufacturing effect (-0.052). Non-metallic minerals exhibit a coefficient of -0.058 (significant at the 1% level), while chemicals show -0.049 (not statistically significant). Importantly, the residual “all other” category—which contains 6,151 observations and is thus not a small sample—shows both a smaller point estimate (-0.028) and is statistically insignificant. This pattern indicates that manufacturing’s aggregate emission response is concentrated in the three major emission-intensive subsectors we isolate, which together account for approximately 65% to 70% of manufacturing emissions and sample observations, rather than being diffused across all manufacturing activities. These sectoral differences likely reflect variation in abatement opportunities across manufacturing subsectors.¹⁹

5.4.3. Subsample stability: Orbis coverage

Our analysis of emission efficiency and real activity outcomes (operating revenue, employment, and assets) requires merging EU ETS emissions data with balance sheet data from Bureau van Dijk’s Orbis database. This reduces the sample size because balance sheet data is not available for all firms reporting emissions in the EU ETS Union Registry. Figure 3 illustrates the sample coverage by presenting annual aggregate emissions for manufacturing and power sectors separately, distinguishing between total EU ETS emissions, emissions from firms in our estimation sample, and emissions from firms with available Orbis data. While such data attrition is inevitable when combining regulatory and commercial databases, it raises an important identification concern: if high-exposed firms differentially exit the Orbis sample after the regulatory tightening—for instance, because regulatory pressure disproport-

tionately affects financially distressed or less transparent firms— our estimated treatment effects could reflect selection bias rather than causal regulatory impact. Even if the composition of Orbis-matched firms remains stable, a related concern is external validity: the emission reductions we document might be specific to firms with formal financial reporting, limiting the generalizability of our findings to the broader EU ETS population.

To assess whether Orbis coverage introduces bias, we test whether treatment effects differ systematically between firms with and without validated financial data. Table OA.6 in the online appendix presents a triple-interaction specification that augments our baseline emissions model with NotOrbis_{ft} , a dummy equal to one for firm-year observations where emissions data exist but cannot be matched to Orbis operating revenue records. The coefficient on $\text{Post}_t \times \text{Exp}_f$ captures the treatment effect for Orbis-matched firms and remains negative and highly significant (-0.122 , $p < 0.01$), consistent with our main results. Critically, the triple interaction $\text{Post}_t \times \text{Exp}_f \times \text{NotOrbis}_{ft}$ tests whether firms without Orbis data exhibit different emission responses. This coefficient is small (-0.042) and statistically insignificant across all specifications and sectors, indicating that emission reductions among high-exposed firms are statistically indistinguishable between Orbis-matched and non-matched subsamples.²⁰ These findings provide strong evidence that the restriction to firms with available financial statements does not introduce selection bias into our (emissions) analyses. The response to the regulatory tightening that we document appears to be a general phenomenon across the EU ETS population rather than an artifact of the Orbis-matched subsample.

5.4.4. *Sample period*

Our sample period starts in 2013 (the beginning of Phase III of the EU ETS) and ends in 2023 (the last year for which data was available at the time of writing). During this period, the EU ETS has seen more than one regulatory tightening: first the announcement of the MSR reform with cancellation-mechanism in 2017, and then the announcement of the

Climate Law in 2020, implemented through the Fit-for-55 package in 2021. In most of our analyses, we have grouped these events together and analyzed their impact jointly through one $Post_t$ dummy which equals 1 over the period 2017-2023, and 0 before. However, as we have shown in Figures 1 and 2, these announcements have each had a significant impact on the carbon price in the EU ETS. It could therefore be interesting to split the post-period into two separate dummies ($Post_t^{2017-2019}$ and $Post_t^{2020-2023}$) and interact them with our exposure dummy (Exp_f) to capture this dimension.

Furthermore, this split is helpful to confirm that our estimates are not confounded by subsequent macroeconomic shocks, most notably the COVID-19 pandemic in 2020 and Russia's invasion of Ukraine in 2022. Note that our difference-in-differences design neutralizes these shocks provided they equally affect high- and low-exposure firms. We think this assumption is highly plausible given that we define carbon price exposure within narrow NACE 4-digit industries. Additionally, our preferred specifications employ industry×year fixed effects, absorbing any unobserved time-varying heterogeneity at this granular level. To completely rule out these confounders, interacting our exposure dummy (Exp_f) with the pre-pandemic indicator ($Post_t^{2017-2019}$) serves as a clean robustness test, demonstrating that the shift in firm behavior pre-dated these later shocks.

The results in the first three columns of Table OA.7 in the Online Appendix show that high-exposed firms reduce their emissions significantly relative to low-exposed firms already during the period 2017-2019, and also confirm that the emissions-reducing effect strengthened substantially during the period 2020-2023. In the full-sample specification (column 1), the difference-in-differences coefficient increases from -0.099^{***} for the period 2017-2019 to -0.193^{***} for the period 2020-2023. This implies that high-exposed firms reduced their emissions by 9.5 percent ($e^{(-0.099)} - 1$) compared to low-exposed firms in the first years after the 2017 regulatory tightening, relative to the pre-period. After 2020, high-exposed firms

reduced their emissions even further to 17.5 percent ($e^{(-0.193)} - 1$) compared to low-exposed firms, relative to the pre-period. The pattern that relative emission reductions of high-exposed firms nearly double during the second part of the tightening period is pronounced in both the power (column 2) and manufacturing (column 3) sectors.

Columns 4 to 6 of Table OA.7 show a similar pattern for emission efficiency. The difference-in-differences coefficient for the full sample (column 4) increases from -0.074^{***} for the period 2017-2019 to -0.110^{***} for the period 2020-2023. Columns 5 and 6 document that this improvement in emission efficiency is persistent during both periods and can be observed in both the power and the manufacturing sector.

Overall, this analysis confirms that our results are not merely caused by other shocks, such as COVID-19 or the Russian invasion in Ukraine, but also that the additional surge in the carbon price after 2020 –which was caused by further regulatory tightening in the EU ETS– led to further emission reductions of high-exposed firms.

5.4.5. *Alternative estimators*

As discussed in Section 4, our preferred empirical strategy estimates the proportional average treatment effect on emission (efficiency) via Poisson QML. An alternative is to estimate equation (4) via OLS even though the dependent variable is log-transformed. This approach has two main drawbacks. First, OLS applied to a log-transformed dependent variable is in general inconsistent (Silva and Tenreyro 2006; Blackburn 2007). Second, firms can reduce their emissions to zero in which case log-transforming the outcome is ill-defined and other commonly used log-like transformations should not be used to approximate percentage changes (Chen and Roth 2024). Nevertheless, Online Appendix Table OA.8 shows that our results are robust to estimating equation (4) by OLS with a log-transformed dependent variable.²¹ Reversely, the M&A analysis, which uses count data can be estimated with Poisson

Quasi-Maximum Likelihood (Correia, Guimarães, and Tom Zylkin 2020) rather than OLS. We show in Table OA.9 that our OLS results on green deal counts are robust to using a two-way fixed effects specifications estimated via PQML.²² Qualitatively, they imply a relative increase of around 214% in the number of green deals completed by high-exposed firms relative to low-exposed firms after the regulatory tightening. This estimate is in line with the descriptive evidence in the right panel of Figure 8: For low-exposed firms, the number of green deals increases on average from 13 per year before 2017 to 21 per year after 2017. In contrast, the number of green deals increases from 16 per year before 2017 to 40 per year after 2017.

5.4.6. Alternative green acquisition classification

In our baseline analysis, we classify M&A deals as ‘green’ using ClimateBERT, a transformer-based language model trained on climate-related text. To validate this approach and assess the robustness of our findings, we develop an alternative rule-based classification algorithm that systematically screens all deal text for environmental content.

Our rule-based classifier extracts text from the deal comments, deal rationale, and editorial fields in Zephyr and searches for climate and environment related terms. The keyword list ranges from highly specific technical terms (e.g., “carbon capture”, “solar farm”, “net zero”) to broader sustainability-related phrases. To minimize false positives, we implement a two-tier matching logic: high-confidence keywords that unambiguously signal environmental intent are accepted without additional validation, while context-dependent terms (e.g., “sustainable”, “green”, “efficiency”) are only counted when environmental context appears within a 100-character window. We also explicitly filter out business jargon where environmental terms are used non-environmentally (e.g., “sustainable growth”, “greenfield site”, “operational efficiency”). Validation against ClimateBERT shows 84.5% overall agreement among the 2,842 deals in our sample. Of the 15.5% disagreement cases, 51.1% are deals that

ClimateBERT flags as green but our stricter criteria do not capture.

We then construct a conservative “intersection” measure that flags a deal as green only if both our keyword-based algorithm and ClimateBERT classify it as environmentally motivated. Using this intersection measure, the qualitative patterns in our M&A analysis remain unchanged: high-exposed firms increase their propensity to engage in green acquisitions after the regulatory tightening relative to low-exposed firms, with effects concentrated in the manufacturing sector and strengthening after 2020 (see Appendix Tables OA.10 and OA.11). While point estimates are somewhat smaller under this more conservative definition, the statistical significance and economic interpretation of our findings are preserved. This confirms that our main conclusions regarding the M&A channel are not artifacts of the ClimateBERT classification method but reflect genuine shifts in acquisition strategies toward environmentally motivated targets.

6. CONCLUSION

We have presented novel evidence on how firms respond to cap-and-trade regulation when carbon prices are sufficiently high. Our analysis delivers three main findings. First, a weak carbon price signal generates little meaningful firm-level response: the pre-2017 EU ETS left emissions largely unchanged among regulated firms. Second, once carbon prices rose sharply following the 2017 regulatory tightening, high-exposed firms reduced emissions while maintaining output, implying efficiency gains rather than output contractions. Third, and most distinctively, the market for corporate control emerges as a channel through which firms realize these efficiency gains: high-exposed manufacturing firms significantly increased their acquisitions of green targets after 2017, while we find no evidence that firms responded by shifting acquisitions toward targets outside the regulated area. The absence of investment leakage through the M&A channel is particularly noteworthy given the prominence of carbon

leakage concerns in the policy debate surrounding the EU ETS.

Our study advances the literature on the firm-level effects of cap-and-trade systems in several respects. Unlike work that exploits early-phase variation in the EU ETS to estimate local average treatment effects (Dechezleprêtre, Nachtigall, and Venmans 2023; Colmer et al. 2025), our difference-in-differences strategy uses the post-2017 regulatory tightening to identify responses across a large, representative sample of regulated firms. The heterogeneity between power producers and manufacturing firms is itself informative: the former achieved large emission reductions consistent with short-run fuel switching, while the latter adjusted through the market for corporate control, acquiring green targets to improve operational efficiency. This distinction matters for policy because it implies that the decarbonization of hard-to-abate manufacturing sectors will depend not only on the price signal but on whether the market for corporate control can sustain green reallocation at scale. A further contribution is the evidence of positive spillovers to unregulated firms: by increasing the attractiveness of green acquisition targets, carbon pricing extends its reach beyond the directly regulated perimeter.

Our findings open two immediate directions for future work. The null result on carbon leakage provides a pre-Carbon Border Adjustment Mechanism (CBAM) benchmark for evaluating subsequent restructuring. CBAM alters incentives for manufacturing firms exposed to foreign competition, creating a natural setting to reassess the geography of green M&A—whether it increases within-Europe acquisitions, redirects deals toward non-EU targets, or attracts inbound green investment. Our estimates provide the relevant counterfactual, particularly for manufacturing, which dominates green M&A in our sample and is CBAM’s primary target. Beyond deal geography, our estimates are identified under historically accommodative macro-financial conditions that shaped both the incentive and ability to restructure. Whether comparable green reallocation occurs under tighter monetary policy,

weaker fiscal support, and greater geopolitical fragmentation is central to assessing carbon pricing as a durable tool of industrial transformation.

LITERATURE CITED

- Abrell, Jan, Mirjam Kosch, and Sebastian Rausch (2022). “How Effective Is Carbon Pricing?—A Machine Learning Approach to Policy Evaluation”. *Journal of Environmental Economics and Management* 112, 102589.
- Bartram, Söhnke M., Kewei Hou, and Sehoon Kim (2022). “Real Effects of Climate Policy: Financial Constraints and Spillovers”. *Journal of Financial Economics* 143 (2), 668–696.
- Bena, Jan and Kai Li (2014). “Corporate innovations and mergers and acquisitions”. *The Journal of Finance* 69 (5), 1923–1960.
- Berg, Tobias, Lin Ma, and Daniel Streitz (2024). *Out of sight, out of mind: Divestments and the global reallocation of pollutive assets*. SAFE Working Paper Series 436. Leibniz Institute for Financial Research SAFE.
- Blackburn, McKinley L. (2007). “Estimating Wage Differentials without Logarithms”. *Labour Economics* 14 (1), 73–98.
- Boeckx, Jana, Kristof Struyfs, and Wouter Torsin (2025). “Green Pressure, Lean Measures: Unveiling Corporate Downsizing Within the European Union Emissions Trading System”. *Journal of Financial and Quantitative Analysis* 60 (8), 4091–4130.
- Bolton, Patrick and Marcin Kacperczyk (2021). “Do investors care about carbon risk?”. *Journal of Financial Economics* 142 (2), 517–549.
- Borghesi, Simone, Michael Pahle, Grischa Perino, Simon Quemin, and Maximilian Willner (2023). “The Market Stability Reserve in the EU Emissions Trading System: A Critical Review”. *Annual Review of Resource Economics* 15 (Volume 15, 2023), 131–152.
- Bruninx, Kenneth, Marten Ovaere, and Erik Delarue (2020). “The Long-Term Impact of the Market Stability Reserve on the EU Emission Trading System”. *Energy Economics* 89, 104746.

- Calel, Raphael and Antoine Dechezleprêtre (2016). “Environmental policy and directed technological change: evidence from the European carbon market”. *Review of economics and statistics* 98 (1), 173–191.
- Cameron, Aliénor and Victor Ho (2024). *Matching the EU Transaction Log and ORBIS: A Natural Language Processing Approach*. https://single-market-economy.ec.europa.eu/single-market/services/economic-analysis/matching-eu-transaction-log-orbis-database_en.
- Chen, Jiafeng and Jonathan Roth (2024). “Logs with Zeros? Some Problems and Solutions*”. *The Quarterly Journal of Economics* 139 (2), 891–936.
- Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J Wagner (2025). “Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from the European Union Emissions Trading System”. *The Review of Economic Studies* 92 (3), 1625–1660.
- Correia, Sergio (2023). “REGHDFE: Stata Module to Perform Linear or Instrumental-Variable Regression Absorbing Any Number of High-Dimensional Fixed Effects”. *Statistical Software Components*.
- Correia, Sergio, Paulo Guimarães, and Thomas Zylkin (2025). *Verifying the Existence of Maximum Likelihood Estimates for Generalized Linear Models*.
- Correia, Sergio, Paulo Guimarães, and Tom Zylkin (2020). “Fast Poisson Estimation with High-Dimensional Fixed Effects”. *The Stata Journal* 20 (1), 95–115.
- Cullen, Joseph A. and Erin T. Mansur (2017). “Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution”. *American Economic Journal: Economic Policy* 9 (3), 106–133.
- De Haas, Ralph, Ralf Martin, Mirabelle Muûls, and Helena Schweiger (2025). “Managerial and Financial Barriers to the Green Transition”. *Management Science* 71 (4), 2890–2921.

- Dechezleprêtre, Antoine, Daniel Nachtigall, and Frank 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.
- Duscha, Vicki, Tobias Fleiter, Matthias Rehfeldt, and Frank Sensfuss (2022). *Abatement Cost Curves for the European Emissions Trading System (EU ETS) Final Report*. Tech. rep. 1862-4359. Germany, p. 247.
- Ellerman, A. Denny, Claudio Marcantonini, and Aleksandar Zaklan (2016). “The European Union Emissions Trading System: Ten Years and Counting”. *Review of Environmental Economics and Policy* 10 (1), 89–107.
- Fell, Harrison and Daniel T. Kaffine (2018). “The Fall of Coal: Joint Impacts of Fuel Prices and Renewables on Generation and Emissions”. *American Economic Journal: Economic Policy* 10 (2), 90–116.
- Fuss, Sabine, Christian Flachsland, Nicolas Koch, Ulrike Kornek, Brigitte Knopf, and Ottmar Edenhofer (2018). “A Framework for Assessing the Performance of Cap-and-Trade Systems: Insights from the European Union Emissions Trading System”. *Review of Environmental Economics and Policy* 12 (2), 220–241.
- Gugler, Klaus, Adhurim Haxhimusa, and Mario Liebensteiner (2023). “Carbon Pricing and Emissions: Causal Effects of Britain’s Carbon Tax”. *Energy Economics* 121, 106655.
- Hanna, Rema (2010). “US Environmental Regulation and FDI: Evidence from a Panel of US-Based Multinational Firms”. *American Economic Journal: Applied Economics* 2 (3), 158–89.
- Harford, Jarrad (2005). “What drives merger waves?” *Journal of Financial Economics* 77 (3), 529–560.
- Ilhan, Emirhan, Zacharias Sautner, and Grigory Vilkov (2021). “Carbon tail risk”. *The Review of Financial Studies* 34 (3), 1540–1571.

- Kalemli-Özcan, Şebnem, Bent E. Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yeşiltaş (2024). “How to Construct Nationally Representative Firm-Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration”. *American Economic Journal: Macroeconomics* 16 (2), 353–374.
- Leroutier, Marion (2022). “Carbon Pricing and Power Sector Decarbonization: Evidence from the UK”. *Journal of Environmental Economics and Management* 111, 102580.
- Martin, Ralf, Mirabelle Muûls, and Ulrich J. Wagner (2016). “The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence after Ten Years?” *Review of Environmental Economics and Policy* 10 (1), 129–148.
- Mitchell, Mark L. and J. Harold Mulherin (1996). “The impact of industry shocks on takeover and restructuring activity”. *Journal of Financial Economics* 41 (2), 193–229.
- Mulier, Klaas, Marten Ovaere, and Leo Stimpfle (2024). *Emission Trading and Overlapping Environmental Support: Installation-level Evidence from the EU ETS*. Working Paper 461. NBB Working Paper.
- Nicholson, Scott and Garvin Heath (2021). “Life Cycle Emissions Factors for Electricity Generation Technologies”. *NLR Data Catalog. Golden, CO: National Laboratory of the Rockies*. Last updated, January 21, 2025.
- Pastor, Lubos, Robert F Stambaugh, and Lucian A Taylor (2022). “Dissecting green returns”. *Journal of Financial Economics* 146 (2), 403–424.
- Perino, Grischa (2018). “New EU ETS Phase 4 Rules Temporarily Puncture Waterbed”. *Nature Climate Change* 8 (4), 262–264.
- Perino, Grischa, Robert A Ritz, and Arthur A van Benthem (2025). “Overlapping Climate Policies”. *The Economic Journal* 135 (671), 2122–2160.

- Perino, Grischa and Maximilian Willner (2016). “Procrastinating Reform: The Impact of the Market Stability Reserve on the EU ETS”. *Journal of Environmental Economics and Management*. The Economics of the European Union Emission Trading System (EU ETS) Market Stability Reserve 80, 37–52.
- Phillips, Gordon M and Alexei Zhdanov (2013). “R&D and the incentives from merger and acquisition activity”. *The Review of Financial Studies* 26 (1), 34–78.
- Rosendahl, Knut Einar (2019). “EU ETS and the Waterbed Effect”. *Nature Climate Change* 9 (10), 734–735.
- Sato, Misato, Ryan Rafaty, Raphael Calel, and Michael Grubb (2022). “Allocation, Allocation, Allocation! The Political Economy of the Development of the European Union Emissions Trading System”. *WIREs Climate Change* 13 (5), e796.
- Silva, J. M. C. Santos and Silvana Tenreiro (2006). “The Log of Gravity”. *The Review of Economics and Statistics* 88 (4), 641–658.
- Sitarz, Joanna, Michael Pahle, Sebastian Osorio, Gunnar Luderer, and Robert Pietzcker (2024). “EU Carbon Prices Signal High Policy Credibility and Farsighted Actors”. *Nature Energy*, 1–12.
- Webersinke, Nicolas, Mathias Kraus, Julia Anna Bingler, and Markus Leippold (2022). *ClimateBert: A Pretrained Language Model for Climate-Related Text*.
- Wooldridge, Jeffrey M (2023). “Simple approaches to nonlinear difference-in-differences with panel data”. *The Econometrics Journal* 26 (3), C31–C66.

FOOTNOTES

1. <https://union-registry-data.ec.europa.eu>

2. This estimate on the full-sample refers to our preferred specification which includes firm, industry-year and country-year fixed effects and is reported in column 4 of Table 2.

3. In the first adopted version of the MSR, allowances in excess of the total number of allowances auctioned in the previous year are canceled (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02015D1814-20180408>) which in 2023 was updated to a fixed value of 400 million allowances (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02015D1814-20240101>).

4. The number of allowances transferred to the MSR equals the total number of allowances in circulation (TNAC) multiplied by the intake rate (initially fixed at 24% conditional on the TNAC exceeding 833 million allowances). The first publication of the TNAC on May 12, 2017 showed a surplus of 1.7 billion allowances and, hence, that hundreds of millions of allowances were likely going to be canceled starting in 2023. It is important to note that while the market appreciated that the cancellation of allowances unambiguously increases regulatory stringency in the EU ETS, the effect of the MSR on aggregate emissions is less clear cut and depends on the timing of emission abatement and banking of allowances (Perino 2018; Rosendahl 2019) as well as overlapping policies (Perino, Ritz, and Benthem 2025).

5. Available at <https://union-registry-data.ec.europa.eu>. We use the pre-processed version provided by Jan Abrell at <https://www.euets.info>.

6. Our sample does not cover all EU ETS emissions because we rely on an installation-level approximation of emission intensity which cannot be applied to all EU ETS installations due to data limitations. See Section 4 for more details.

7. In the robustness section, we check whether this matching step induces a selection bias that potentially

impacts the results and show that it does not.

8. We use the downstream classification task `climatebert/distilroberta-base-climate-detector`

9. <https://globalenergymonitor.org/projects/global-integrated-power-tracker/>

10. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32011D0278>

11. We refer to (Mulier, Ovaere, and Stimpfle 2024) for further details.

12. In the robustness section (5.4.5), we elaborate on this choice and show that using alternative estimators yield qualitatively similar results both for the emission (efficiency) and M&A analyses.

13. For example, Dechezleprêtre, Nachtigall, and Venmans (2023) restrict their sample period to 2000–2012 “because the coverage of the EU ETS broadened substantially in 2013 with the beginning of the third trading phase, leaving only few control observations in some sectors.”

14. Our analysis of emission efficiency and real activity outcomes requires merging EU ETS emissions data with balance sheet data from Bureau van Dijk’s Orbis database. This reduces the sample size because balance sheet data is not available for all firms reporting emissions in the EU ETS Union Registry. While such data attrition is inevitable when combining regulatory and commercial databases, it raises potential concerns about selection bias if high-exposed firms differentially exit the Orbis sample after the regulatory tightening, or about external validity if our findings are specific to firms with formal financial reporting. In Section 5.4.3, we demonstrate that effects on emissions shown above are statistically indistinguishable between Orbis-matched and non-matched firms, and that high-exposed firms do not differentially drop out of Orbis coverage after 2017. These tests provide strong evidence that the restriction to firms with available financial statements does not introduce bias into our emission efficiency analysis.

15. We also obtain null effects on total assets (showing positive and statistically insignificant coefficients) and employment (showing negative and statistically insignificant coefficients, except for manufacturing but

additional results show this is driven by a quirk in 2013 and 2014), which we report in Online Appendix Table OA.1.

16. In the robustness section 5.4.6, we confirm these results using an alternative text-based classification.

17. Five percentage points for high-exposed power firms is the difference between $e^{-0.391}$ and $e^{-0.321}$ and one percentage point for high-exposed manufacturing firms is the difference between $e^{-0.052}$ and $e^{-0.043}$.

18. We note that before the regulatory tightening high-exposed firms were more likely to acquire targets not located in EU ETS member states. This, however, is likely driven by high-exposed firms' higher propensity to engage in M&As (see Table 4).

19. Basic metals producers, particularly steel manufacturers, have several adjustment margins including fuel switching, energy efficiency improvements, and strategic acquisitions to increase scrap metal usage. The latter being illustrated by our anecdotal M&A evidence in Online Appendix A, where ArcelorMittal's acquisition of a scrap processing facility was explicitly motivated by emission reduction objectives.

20. We also verify that the probability of having Orbis data available is not systematically correlated with treatment exposure and timing. Estimating a linear probability model with an Orbis match indicator as the dependent variable, we find that the interaction $\text{Post}_t \times \text{Exp}_f$ is close to zero (0.6% in our preferred specification) and statistically insignificant, confirming that high-exposure firms do not differentially drop out of Orbis after the regulatory tightening. Full results are available upon request.

21. Ignoring that estimates of the latter are potentially biased (Blackburn 2007; Silva and Tenreyro 2006), the difference in coefficient magnitude can also be related to a nuance in their respective interpretations: Whereas Poisson QML estimates the ATE in levels relative to the baseline mean, OLS applied to the log-transformed dependent variable estimates the ATE of the log-difference in the outcome (Chen and Roth 2024). Consequently, Poisson QML is less likely to be affected by extreme relative changes of small emitters and better reflects the evolution of aggregate emission changes in the EU ETS. See also the instructive discussion in Annex F of (Dechezleprêtre, Nachtigall, and Venmans 2023).

22. We note that the sample size drops materially to 955 observations covering 133 firms. This is because the STATA command `ppmlhdfc` identifies and drops so-called “separated” observations to avoid issues related to the nonexistence of quasi-maximum likelihood estimates (Correia, Guimarães, and Thomas Zylkin 2025, Proposition 1). As a consequence, firms that do not engage in any deals are dropped from the estimation sample in Table OA.9.

TABLES

Table 1: Summary statistics

Sample	Variable	Mean	Standard Deviation	5%	25%	50%	75%	95%	Observations
Full	Installations	2.3	3.6	1	1	1	2	8	24,202
	Emissions (tCO ₂ e)	558,374.1	2,820,569.9	140	10,561	44,316	203,352	2,467,187	24,202
	Operating Revenue (mEUR)	900.0	8,134.7	4	33	98	290	2,022	15,848
	M&As	0.1	0.7	0	0	0	0	1	20,269
	Green M&As	0.0	0.2	0	0	0	0	0	20,269
Power	Installations	4.3	5.7	1	1	2	6	14	3,738
	Emissions (tCO ₂ e)	1,957,324.9	6,575,199.5	754	83,146	337,824	1,291,829	7,646,021	3,738
	Operating Revenue (mEUR)	2,390.9	18,064.2	19	63	187	612	3,744	2,443
	M&As	0.3	1.1	0	0	0	0	2	2,819
	Green M&As	0.1	0.5	0	0	0	0	1	2,819
Manufacturing	Installations	2.0	2.9	1	1	1	2	5	20,430
	Emissions (tCO ₂ e)	301,753.7	1,043,508.9	36	9,379	34,025	125,423	1,600,678	20,430
	Operating Revenue (mEUR)	625.2	4,287.1	3	28	88	255	1,512	13,338
	M&As	0.1	0.6	0	0	0	0	1	17,405
	Green M&As	0.0	0.2	0	0	0	0	0	17,405

Notes: The table presents summary statistics of the main variables. The sample period is 2013–2023 and each observation represents a firm-year.

Table 2: Emissions and EU ETS regulatory tightening

Dependent Sample Specification	Emissions					
	(1)	(2)	Full (3)	(4)	Power (5)	Manufacturing (6)
$Post_t$	-0.041 (0.049)					
Exp_f	1.306*** (0.171)					
$Post_t \times Exp_f$	-0.214*** (0.057)	-0.210*** (0.057)	-0.110*** (0.031)	-0.139*** (0.027)	-0.391*** (0.061)	-0.052*** (0.016)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	24,202	24,202	24,202	24,202	3,738	20,430
Firms	2,207	2,207	2,207	2,207	343	1,861
Pseudo R^2	0.075	0.982	0.985	0.986	0.975	0.994

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm’s emissions (in ton CO₂ equivalents). $Post_t$ is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3: Emission efficiency and EU ETS regulatory tightening

Dependent Sample Specification	Operating Revenue			Emissions / Operating Revenue		
	Full (1)	Power (2)	Manufacturing (3)	Full (4)	Power (5)	Manufacturing (6)
$\text{Post}_t \times \text{Exp}_f$	-0.025 (0.031)	-0.054 (0.062)	-0.008 (0.029)	-0.089*** (0.023)	-0.239*** (0.058)	-0.056*** (0.021)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Fuel price controls				Yes	Yes	
Observations	15,848	2,443	13,338	15,848	2,443	13,338
Firms	1,619	245	1,368	1,619	245	1,368
Pseudo R^2	0.997	0.997	0.997	0.618	0.620	0.594

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm’s operating revenue (in thousand euro) in columns 1 to 3 and firms’ emission (in)efficiency, proxied by the ratio of total emissions (in tonne CO₂ equivalent) over operating revenue (in thousand euro) in columns 4 to 6. Post_t is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: M&A deals

Dependent Sample Specification	Deals					
	(1)	Full (2)	(3)	(4)	Power (5)	Manufacturing (6)
$Post_t$	0.003 (0.009)					
Exp_f	0.055** (0.024)					
$Post_t \times Exp_f$	0.004 (0.017)	0.004 (0.017)	0.012 (0.015)	0.008 (0.016)	-0.120 (0.093)	0.011 (0.016)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	20,269	20,269	20,269	20,269	2,819	17,405
Firms	1,847	1,847	1,847	1,847	257	1,586
R^2	0.002	0.628	0.629	0.644	0.681	0.641

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the number of M&A deals that a firm announces in a given year. $Post_t$ is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm's exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. All specifications are estimated via OLS. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 5: Green M&A deals

Dependent Sample Specification	Green Deals					
	(1)	(2)	Full (3)	(4)	Power (5)	Manufacturing (6)
Post_t	0.009** (0.003)					
Exp_f	0.002 (0.005)					
$\text{Post}_t \times \text{Exp}_f$	0.016** (0.007)	0.015** (0.007)	0.015** (0.007)	0.017** (0.008)	-0.018 (0.060)	0.019** (0.008)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	20,269	20,269	20,269	20,269	2,819	17,405
Firms	1,847	1,847	1,847	1,847	257	1,586
R^2	0.002	0.487	0.488	0.509	0.542	0.509

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the number of green M&A deals that a firm announces in a given year. We use the pretrained large language model ClimateBERT (Webersinke et al. 2022) on the text incorporated as “deal rationale” in ZEPHYR to classify an M&A as “green” or not. Post_t is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated via OLS. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Green acquisitions and emission reductions

Dependent Sample Specification	Emissions			
	Full (1)	Full (2)	Orbis (3)	Orbis (4)
$M\&A_{f,t-1}^{\text{green}}$	-0.119** (0.052)	-0.098* (0.052)	-0.118** (0.050)	-0.105** (0.044)
$M\&A_{f,t-1}^{\text{green}} \times \text{Post}_t$				-0.019 (0.040)
$\ln(\text{OperatingRevenue}_{ft})$			0.081** (0.032)	0.078*** (0.030)
$\ln(\text{TotalAssets}_{ft})$			-0.025 (0.031)	-0.024 (0.030)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes			
Year-country FE		Yes	Yes	Yes
Year-industry FE		Yes	Yes	Yes
Fuel price controls	Yes	Yes	Yes	Yes
Observations	20,269	20,269	12,849	12,849
Firms	1,847	1,847	1,311	1,311
Pseudo R^2	0.989	0.990	0.993	0.993

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm’s emissions (in ton CO₂ equivalents). $M\&A_{f,t-1}^{\text{green}}$ is an indicator equal to one if firm f has completed at least one green deal by year $t - 1$, and zero otherwise. Green deals are identified using the ClimateBERT classification. Post_t is a dummy equal to one if $t \geq 2017$, and zero otherwise. Columns 1 and 2 present results on the full sample; columns 3 and 4 present results on a restricted sample for which log operating revenue and log total assets are available as controls. All specifications include firm fixed effects. Column 1 includes year fixed effects; columns 2–4 replace year fixed effects with year-country and year-industry fixed effects. All specifications include fuel price controls for power producers. Columns 3 and 4 additionally include log operating revenue and log total assets as controls. Column 4 adds an interaction between the green M&A indicator and the post-period dummy. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Financial constraints

Dependent Sample Specification	Green Deals								
	Full (1)	Power (2)	Manufacturing (3)	Full (4)	Power (5)	Manufacturing (6)	Full (7)	Power (8)	Manufacturing (9)
$\text{Post}_t \times \text{Exp}_f$	0.024* (0.014)	-0.031 (0.096)	0.021 (0.014)	0.022** (0.011)	0.045 (0.117)	0.016* (0.010)	0.029* (0.016)	-0.053 (0.123)	0.030* (0.016)
$\text{Post}_t \times \mathbb{1}(\text{Cash}_f)$	-0.005 (0.012)	0.080 (0.097)	-0.012 (0.010)						
$\text{Post}_t \times \text{Exp}_f \times \mathbb{1}(\text{Cash}_f)$	0.001 (0.020)	-0.065 (0.164)	0.006 (0.016)						
$\text{Post}_t \times \mathbb{1}(\text{Leverage}_f)$				0.002 (0.011)	0.092 (0.075)	-0.005 (0.009)			
$\text{Post}_t \times \text{Exp}_f \times \mathbb{1}(\text{Leverage}_f)$				0.007 (0.023)	-0.194 (0.151)	0.022 (0.022)			
$\text{Post}_t \times \mathbb{1}(\text{Size}_f)$							0.004 (0.010)	-0.119 (0.151)	0.010 (0.010)
$\text{Post}_t \times \text{Exp}_f \times \mathbb{1}(\text{Size}_f)$							-0.016 (0.017)	0.146 (0.171)	-0.020 (0.017)
Fuel price controls	Yes	Yes		Yes	Yes		Yes	Yes	
Observations	13,902	1,232	13,146	13,902	1,232	13,146	13,902	1,232	13,146
Firms	1,266	112	1,197	1,266	112	1,197	1,266	112	1,197
R^2	0.528	0.578	0.531	0.528	0.578	0.531	0.528	0.578	0.531

Notes: The sample consists of firm-year level observations for the period 2013–2023, restricted to firms with validated Orbis balance sheet data available. This specification tests whether firm financial characteristics moderate the response to regulatory tightening in terms of green M&A activity. The dependent variable is the number of green M&A deals that a firm announces in a given year. We use the pretrained large language model ClimateBERT (Webersinke et al. 2022) on the text incorporated as “deal rationale” in ZEPHYR to classify an M&A as “green” or not. Post_t is a dummy equal to one if $t \geq 2017$, and zero otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. Columns 1–3 test heterogeneity by cash holdings: $\mathbb{1}(\text{Cash}_f)$ is a dummy equal to one for firms in the bottom tercile of the cash-to-assets ratio distribution over 2013–2016 (i.e., firms with low cash holdings), and zero otherwise. Columns 4–6 test heterogeneity by leverage: $\mathbb{1}(\text{Leverage}_f)$ is a dummy equal to one for firms in the top tercile of the leverage ratio distribution over 2013–2016 (i.e., highly leveraged firms), and zero otherwise. Columns 7–9 test heterogeneity by firm size: $\mathbb{1}(\text{Size}_f)$ is a dummy equal to one for firms in the bottom tercile of the total assets distribution over 2013–2016 (i.e., small firms), and zero otherwise. A firm’s exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. The coefficients of interest are the triple interactions: a significant coefficient would indicate that the corresponding firm characteristic moderates the treatment effect. All specifications include firm fixed effects, year-country fixed effects, and year-industry fixed effects. Fuel price controls are included for the full sample and power producers (columns 1–2, 4–5, and 7–8). Depending on the fixed effects included, some level terms and double interactions are absorbed and not separately estimated. All specifications are estimated via OLS. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 8: M&A deals with targets outside of EU ETS member states

Dependent Sample Specification	Target outside EU ETS zone					
	(1)	Full	(3)	(4)	Power	Manufacturing
		(2)			(5)	(6)
$Post_t$	0.003 (0.003)					
Exp_f	0.015** (0.007)					
$Post_t \times Exp_f$	-0.001 (0.006)	-0.001 (0.006)	-0.000 (0.006)	-0.001 (0.006)	-0.031 (0.035)	-0.002 (0.007)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	20,269	20,269	20,269	20,269	2,819	17,405
Firms	1,847	1,847	1,847	1,847	257	1,586
R^2	0.001	0.496	0.496	0.515	0.548	0.518

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the number of M&A deals with a target located outside of EU ETS member states that a firm announces in a given year. $Post_t$ is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm's exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. All specifications are estimated via OLS. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.


FIGURES



Figure 1: EU ETS market prices and the MSR reform


Notes: The figure presents the price evolution of EUAs (EU ETS emissions allowances) since the start of EU ETS Phase III in 2013. The price data is based on the front December future contract and obtained from Refinitiv (ticker CFI2Zc1). The vertical dashed lines indicate key regulatory events in the context of the Market Stability Reserve (MSR) reform. The MSR was first proposed by the EU Commission on July 16, 2015 (https://eur-lex.europa.eu/procedure/EN/2015_148#2017-02-15_OPI_R1_byEP) without substantively affecting EUA prices. On 15 February 2017 the European Parliament published its opinion on the first reading of the EU ETS reform which significantly strengthened the EU Commission's proposal, most notably by a mechanism to permanently cancel excess allowances in the MSR (https://www.europarl.europa.eu/doceo/document/TA-8-2017-0035_EN.html). After the first publication of the Total Number of Allowances in Circulation (TNAC) on May 12, 2017 ([https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52017XC0513\(01\)](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52017XC0513(01))), market participants could for the first time accurately estimate the number of allowances to be transferred to the MSR and the effect of cancellations on the future supply of allowances triggering an unprecedented rally in EUA prices.

Figure 2: EU ETS market prices and key regulatory events




Notes: The figure presents the price evolution of EUAs (EU ETS emissions allowances) since the start of EU ETS Phase III in 2013. The price data is based on the front December future contract and obtained from Refinitiv (ticker CFI2Zc1). The vertical dashed lines indicate key regulatory events since EU ETS Phase III. For a description of events relating to the Market Stability Reserve reform (“Cancellation” and “1st TNAC”) see Figure 1. After the MSR reform triggered a rally leading to a quadrupling of EUA prices, prices stabilized around 25 EUR/tCO₂ over 2019 and 2020. Signs of EU’s increasingly ambitious 2030 decarbonization targets first transpired in the proposal of the Climate Law on March 4, 2020 (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2020:0080:FIN>) which were later implemented as a materially lower annual EU ETS cap and a higher MSR intake in the so-called “Fit-for-55” package (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0551>) driving prices to new highs of up to 100 EUR/tCO₂ in February 2023.

Figure 3: Sample coverage of emissions (by broad EU ETS sectors)



Notes: The figure presents the annual aggregate emissions of firms in the power sample (left panel) and the manufacturing sample (right panel). For each panel we show the breakdown of emissions of (i) firms in the EU ETS, (ii) firms in the EU ETS with installations for which we can credibly approximate the emission intensity, and (iii) firms in the EU ETS with installations for which we can credibly approximate the emission intensity and match to Orbis with validated balance sheet data following Kalemli-Özcan et al. (2024). A firm's emission intensity is defined as the emission-weighted intensity of its installations. An installation's emission intensity is inferred from the Global Energy Monitor (power) or the harmonized free allocation methodology (manufacturing).

Figure 4: DiD event study on emissions: full sample




Notes: The figure presents Poisson QML estimates (Wooldridge 2023) of the event study specification

$$\text{Emissions}_{f,t} = \sum_{s=2013}^{2023} \beta_s \times \mathbb{1}_s(t) \times \text{Exp}_f + \alpha_f + X_{f,t} + \gamma_{c,t} + \delta_{i,t} + \epsilon_{f,t},$$

where $\mathbb{1}_s(t)$ is a dummy equal to one when year t equals s , α_f are firm fixed effects, $\gamma_{c,t}$ are country-year fixed effects, $\delta_{i,t}$ are industry-year fixed effects, and Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. To account for time-varying marginal abatement costs of power producers, we control for fuel prices (Cullen and Mansur 2017) and the annual country-level share of renewable energy (Fell and Kaffine 2018). A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. The event study coefficients β_s are estimated relative to 2016—the last year before announcement of the regulatory tightening—using the STATA command `ppmlhdfc` (Correia, Guimarães, and Tom Zylkin 2020). Standard errors are clustered at the firm level and 90% confidence intervals are shown.

Figure 5: DiD event study on emissions




Notes: The figure presents Poisson QML estimates (Wooldridge 2023) of the event study specification

$$\text{Emissions}_{f,t} = \sum_{s=2013}^{2023} \beta_s \times \mathbb{1}_s(t) \times \text{Exp}_f + \alpha_f + X_{f,t} + \gamma_{c,t} + \delta_{i,t} + \epsilon_{f,t},$$

where $\mathbb{1}_s(t)$ is a dummy equal to one when year t equals s , α_f are firm fixed effects, $\gamma_{c,t}$ are country-year fixed effects, $\delta_{i,t}$ are industry-year fixed effects, and Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. Panel (a) presents results for power producers. Panel (b) presents results for manufacturing firms. To account for time-varying marginal abatement costs of power producers, we control for fuel prices (Cullen and Mansur 2017) and the annual country-level share of renewable energy (Fell and Kaffine 2018) in panel (a). A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. The event study coefficients β_s are estimated relative to 2016—the last year before announcement of the regulatory tightening—using the STATA command `ppmlhdfc` (Correia, Guimarães, and Tom Zylkin 2020). Standard errors are clustered at the firm level and 90% confidence intervals are shown.

Figure 6: DiD event study on emissions over operating revenue: full sample




Notes: The figure presents Poisson QML estimates (Wooldridge 2023) of the event study specification

$$\text{Emissions/Operating Revenue}_{f,t} = \sum_{s=2013}^{2023} \beta_s \times \mathbb{1}_s(t) \times \text{Exp}_f + \alpha_f + X_{f,t} + \gamma_{c,t} + \delta_{i,t} + \epsilon_{f,t},$$

where $\mathbb{1}_s(t)$ is a dummy equal to one when year t equals s , α_f are firm fixed effects, $\gamma_{c,t}$ are country-year fixed effects, $\delta_{i,t}$ are industry-year fixed effects, and Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. The dependent variable is the ratio of firm emissions (in ton CO₂ equivalents) to operating revenue (in thousand euros), measuring emission efficiency. To account for time-varying marginal abatement costs of power producers, we control for fuel prices (Cullen and Mansur 2017) and the annual country-level share of renewable energy (Fell and Kaffine 2018). A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. The event study coefficients β_s are estimated relative to 2016—the last year before announcement of the regulatory tightening—using the STATA command `ppmlhdfc` (Correia, Guimarães, and Tom Zylkin 2020). Standard errors are clustered at the firm level and 90% confidence intervals are shown.

Figure 7: DiD event study on emissions over operating revenue




Notes: The figure presents Poisson QML estimates (Wooldridge 2023) of the event study specification

$$\text{Emissions/Operating Revenue}_{f,t} = \sum_{s=2013}^{2023} \beta_s \times \mathbb{1}_s(t) \times \text{Exp}_f + \alpha_f + X_{f,t} + \gamma_{c,t} + \delta_{i,t} + \epsilon_{f,t},$$


where $\mathbb{1}_s(t)$ is a dummy equal to one when year t equals s , α_f are firm fixed effects, $\gamma_{c,t}$ are country-year fixed effects, $\delta_{i,t}$ are industry-year fixed effects, and Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. The dependent variable is the ratio of firm emissions (in ton CO₂ equivalents) to operating revenue (in thousand euros), measuring emission efficiency. Panel (a) presents results for power producers. Panel (b) presents results for manufacturing firms. To account for time-varying marginal abatement costs of power producers, we control for fuel prices (Cullen and Mansur 2017) and the annual country-level share of renewable energy (Fell and Kaffine 2018) in panel (a). A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. The event study coefficients β_s are estimated relative to 2016—the last year before announcement of the regulatory tightening—using the STATA command `ppmlhdf` (Correia, Guimarães, and Tom Zylkin 2020). Standard errors are clustered at the firm level and 90% confidence intervals are shown.

Figure 8: Number of (green) M&A deals (conditional on firms' carbon price exposure)



Notes: The figure presents the number of merger and acquisitions that are being announced by firms in our sample during the sample period, conditional on whether a firm's exposure to the carbon price shock caused by the regulatory tightening is relatively high ($\text{Exp}_f = 1$) or relatively low ($\text{Exp}_f = 0$). Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. The left panel contains all deals, whereas the right panel counts the green M&A deals. A firm's exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). For power producers, the emission intensity of an installation is inferred from the Global Energy Monitor. For manufacturing firms, the emission intensity of an installation is measured as the verified emissions of that installation scaled by the tonnes of good produced at that installation (which we infer from the harmonized free allocation methodology). When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions.

Figure 9: DiD event study: green M&A deals



Notes: The figure presents estimates of the event study specification

$$\text{Green M\&A Deals}_{f,t} = \sum_{s=2013}^{2023} \beta_s \times \mathbb{1}_s(t) \times \text{Exp}_f + \alpha_f + \gamma_{c,t} + \delta_{i,t} + \epsilon_{f,t},$$

where Green M&A Deals_{*f,t*} is the number of Green M&A deals announced by firm *f* in year *t*, $\mathbb{1}_s(t)$ is a dummy equal to one when year *t* equals *s*, α_f are firm fixed effects, $\gamma_{c,t}$ are country-year fixed effects, $\delta_{i,t}$ are industry-year fixed effect, and Exp_{*f*} is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. Panel (a) presents results for power producers. Panel (b) presents results for manufacturing firms. A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. The event study coefficients β_s are estimated relative to 2016 –the last year before the announcement of the regulatory tightening– using the STATA command `reghdfe` (Correia 2023). Standard errors are clustered at the firm level and confidence intervals are shown at the 90%-level.

ONLINE APPENDIX

This Online Appendix accompanies the paper “**Climate Regulation, Firm Emissions, and Green Takeovers**”, by Olivier De Jonghe, Klaas Mulier, Glenn Schepens, and Leonard Stimpfle.

This Online Appendix provides additional details on various aspects:

- Section A gives more insight into the classification of deals by ClimateBERT and some examples of green and non-green deals.
- Section B contains additional figures
 - Figure OA.1: Share of global emissions covered by emission trading systems (ETS) and number of active ETS
 - Figure OA.2: Average emissions, conditional on firms’ carbon price exposure (power)
 - Figure OA.3: Average emissions, conditional on firms’ carbon price exposure (manufacturing)
- Section C contains additional tables
 - Table OA.1: Real Activity Outcomes: operating revenue, Total Assets, and Employment
 - Table OA.2: Impact of EU ETS regulatory tightening on CO₂ emissions: with and without green M&A deals
 - Table OA.3: Standardised continuous treatment
 - Table OA.4: Treatment quartiles
 - Table OA.5: Heterogeneity by NACE 2-digit industry
 - Table OA.6: Testing for Differential Treatment Effects: Orbis vs. Non-Orbis

Firms

- Table OA.7: Separate post dummies for 2017–2019 and 2020–2023
- Table OA.8: Alternative estimators
- Table OA.9: Green deals with PQML
- Table OA.10: Number of green deals (intersection of ClimateBERT and alternative definition)
- Table OA.11: Number of green deals - subperiods (intersection of ClimateBERT and alternative definition)

A. TEXTUAL ANALYSIS AND CLASSIFICATION OF ‘GREEN’ M&A DEALS

We first construct the variable $\text{Green M\&A}_{f,d,t}$ based on textual information about merger and acquisition deals. The text is taken from the variable called ‘deal rationale’ in the Zephyr database. Using this text on the pretrained large language model (LLM) ClimateBERT (Webersinke et al. 2022) classifies a deal d as either ‘green’ ($\text{Green M\&A}_{f,d,t} = 1$) or not ($\text{Green M\&A}_{f,d,t} = 0$).

From this, we then calculate the variable $\text{Green M\&A Deals}_{f,t}$ as the number of M&As that firm f announced in year t and that can be considered to have a climate-related (i.e., green) goal. Basically, we take the sum of $\text{Green M\&A}_{f,d,t}$ over all deals done by firm f in year t .

We provide six examples of deal rationales of which three are classified green and three not green by ClimateBERT.

Example 1 (green)

On 04/06/19 Mr Frank Mastiaux, EnBW CEO, said: “The acquisition of VALECO marks a significant step forward in the rigorous expansion of EnBW in renewable energy to make them one of the main pillars of the company. In addition, the target of reaching 1,000 MW of installed capacity in the onshore wind sector by 2020 has now nearly been achieved. With VALECO, we now have one of the most experienced players on the French renewable energy market at our side. We will exploit the growth opportunities together and become one of the Top 5 players on the French wind and solar market in the medium term as strong partners” On 28/03/19 Mr Frank Mastiaux, EnBW CEO, commented: ‘For EnBW, this is another important step on the way to making renewable energy a cornerstone of the

company's business. We see good prospects for common growth in France.' (Deal number 1941246267)

Example 2 (green)

On 07/06/17 Mr Tomas Pleskac, CEZ Group chief renewable energy and distribution officer, stated: "This is our first project in the French renewable energy market, and we believe that more will follow. With similar acquisitions in countries with stable regulatory environments, we continue to fulfil our strategy aimed at increasing our operating profit before depreciation from renewable sources by CZK 3 bn (EUR 114 mil) by 2020. We are interested in completed onshore projects as well as in offshore projects and projects under development." (Deal number 1909589305)

Example 3 (green)

On 02/03/22 Mr Geert Van Poelvoorde, CEO ArcelorMittal Europe, said: "We have identified strong potential for growth in the ferrous scrap processing business, with demand growth in Europe facilitated by the European Union's initiatives to achieve higher metal recycling rates, reduce CO2 emissions and underpin the EU's net-zero ambitions. We are therefore very pleased to announce the acquisition of John Lawrie Metals, which represents a further step in our strategy to increase the use of scrap steel across our steelmaking sites" (Deal number 1941661433)

Example 4 (not green)

On 08/12/14 Mr Raimar Jahn, President, Performance Materials, BASF, said: "The acquisition is part of our efforts to systematically pursue growth in the attractive TPU market. It is a valuable addition to our existing TPU portfolio; strengthening our manufacturing footprint in Asia. As TWSS has been at the forefront of TPU adhesives innovation, the acquisition is a strategic move that will strengthen BASF's competences in this important

growth field.” (Deal number 1907080783)

Example 5 (not green)


On 26/02/14 Mr Lakshmi N. Mittal, chairman and CEO of ArcelorMittal, stated: “The successful completion of this transaction is an important milestone for ArcelorMittal. Along with NSSMC we are now the owners of the most modern steel finishing facility in the world, which will allow us to meet rising demand for steels in the automotive, energy and other important NAFTA markets.” (Deal number 1601337266)

Example 6 (not green)

On 04/02/19 Ms Jan Jenisch, CEO: “In line with our Strategy 2022 – ‘Building for Growth’, these acquisitions will generate synergies with LafargeHolcim’s existing operations. With these further bolt-on acquisitions we are delivering on our commitment to accelerate growth in the Ready-Mix Concrete and Aggregates segments. I am pleased to welcome all new employees to LafargeHolcim.” (Deal number 1941191332)

B. SUPPLEMENTARY FIGURES

Figure OA.1: Share of global emissions covered by emission trading systems (ETS) and number of active ETS




Notes: This figure presents the share of global emissions covered by emissions trading systems (left panel) and the number of active emissions trading systems around the world (right panel). Source: World Bank Carbon Pricing Dashboard.

Figure OA.2: Average emissions, conditional on firms' carbon price exposure (power)



Notes: The figure presents the average emissions expressed in tCO₂ by firms in our sample during the sample period, conditional on whether a firm's exposure to the carbon price shock caused by the regulatory tightening is relatively high ($\text{Exp}_f = 1$) or relatively low ($\text{Exp}_f = 0$). Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. A firm's exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. If a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's average verified emissions in the period 2013–2016. The left panel shows the absolute average over time, and the right panel shows the evolution normalized to one in 2016.

Figure OA.3: Average emissions, conditional on firms' carbon price exposure (manufacturing)



Notes: The figure presents the average emissions expressed in tCO₂ by firms in our sample during the sample period, conditional on whether a firm's exposure to the carbon price shock caused by the regulatory tightening is relatively high ($Exp_f = 1$) or relatively low ($Exp_f = 0$). Exp_f is a dummy equal to one if a firm's carbon price exposure is above the median in its NACE 4-digit industry. A firm's exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. If a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's average verified emissions in the period 2013–2016. The left panel shows the absolute average over time, and the right panel shows the evolution normalized to one in 2016.

C. SUPPLEMENTARY TABLES

Table OA.1: Real Activity Outcomes: operating revenue, Total Assets, and Employment

Dependent Sample Specification	Operating Revenue			Total Assets			Employees		
	Full (1)	Power (2)	Manufacturing (3)	Full (4)	Power (5)	Manufacturing (6)	Full (7)	Power (8)	Manufacturing (9)
$\text{Post}_t \times \text{Exp}_f$	-0.025 (0.031)	-0.054 (0.062)	-0.008 (0.029)	0.015 (0.033)	0.042 (0.068)	0.035 (0.040)	-0.044 (0.028)	-0.058 (0.059)	-0.049* (0.027)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15,848	2,443	13,338	15,788	2,428	13,294	14,301	2,045	12,173
Firms	1,619	245	1,368	1,617	245	1,366	1,573	223	1,343
Pseudo R^2	0.997	0.997	0.997	0.998	0.999	0.996	0.997	0.999	0.989

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variables are operating revenue (columns 1–3), total assets (columns 4–6), and number of employees (columns 7–9), all measured at the firm level. Post_t is a dummy equal to one if $t \geq 2017$, and zero otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. All specifications include firm fixed effects, year-country fixed effects, and year-industry fixed effects. Fuel price controls are included in specifications for the full sample and power producers (columns 1–2 and 4–5). All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.2: Impact of EU ETS regulatory tightening on CO₂ emissions: with and without green M&A deals

Sample	Full		Power		Manufacturing	
Description	Full sample	Excluding green acquirors	Full sample	Excluding green acquirors	Full sample	Excluding green acquirors
$\text{Post}_t \times \text{Exp}_f$	-0.139*** (0.027)	-0.111*** (0.028)	-0.391*** (0.061)	-0.321*** (0.072)	-0.052*** (0.016)	-0.043** (0.018)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE						
Year-country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Fuel price controls	Yes	Yes	Yes	Yes		
Observations	24,202	21,274	3,738	2,646	20,430	18,572
Firms	2,207	1,940	343	243	1,861	1,692
Pseudo R^2	0.986	0.989	0.975	0.982	0.994	0.993

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm’s emissions (in ton CO₂ equivalents). Post_t is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. We use the pretrained large language model ClimateBERT (Webersinke et al. 2022) on the text incorporated as “deal rationale” in ZEPHYR to classify an M&A as “green” or not. If a firm does one green M&A during the sample period, it is excluded from the analysis in columns 2, 4, and 6. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.3: Standardised continuous treatment

Dependent Sample Specification	Emissions					
	(1)	Full (2)	(3)	(4)	Power (5)	Manufacturing (6)
Post_t	-0.167*** (0.023)					
Intensity_f	0.507*** (0.057)					
$\text{Post}_t \times \text{Intensity}_f$	-0.049*** (0.013)	-0.130*** (0.026)	-0.061*** (0.017)	-0.118*** (0.026)	-0.200*** (0.060)	-0.050*** (0.011)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	24,202	24,202	24,202	24,202	3,738	20,430
Firms	2,207	2,207	2,207	2,207	343	1,861
Pseudo R^2	0.100	0.982	0.985	0.986	0.975	0.994

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm’s emissions (in ton CO₂ equivalents). Post_t is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Intensity_f is our emission intensity measure, defined as the emission-weighted average emission intensity of its installations over the period 2013–2016 (see Section 4.1 Difference-in-differences setup in the paper for more details). For ease of interpretation, this variable is standardized to have zero mean and a standard deviation of one. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.4: Treatment quartiles

Dependent Sample Specification	Emissions					
	(1)	Full (2)	(3)	(4)	Power (5)	Manufacturing (6)
$Post_t$	0.050 (0.044)					
$IntensityQ2_f$	0.535*** (0.197)					
$IntensityQ3_f$	1.209*** (0.171)					
$IntensityQ4_f$	1.770*** (0.238)					
$Post_t \times IntensityQ2_f$	-0.176** (0.084)	-0.172** (0.083)	-0.128* (0.070)	-0.130* (0.075)	-0.258 (0.157)	-0.031 (0.023)
$Post_t \times IntensityQ3_f$	-0.348*** (0.069)	-0.339*** (0.068)	-0.188*** (0.054)	-0.204*** (0.054)	-0.520*** (0.133)	-0.049** (0.022)
$Post_t \times IntensityQ4_f$	-0.277*** (0.056)	-0.273*** (0.055)	-0.185*** (0.048)	-0.235*** (0.056)	-0.548*** (0.135)	-0.098*** (0.025)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	24,202	24,202	24,202	24,202	3,738	20,430
Firms	2,207	2,207	2,207	2,207	343	1,861
Pseudo R^2	0.086	0.982	0.985	0.986	0.975	0.994

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm's emissions (in ton CO₂ equivalents). $Post_t$ is a dummy equal to one if $t \geq 2017$, and 0 otherwise. $IntensityQ2_f$ is a dummy equal to one for all firms with carbon intensity in the second quartile in their NACE 4-digit industry. Similarly, $IntensityQ3_f$ and $IntensityQ4_f$ are dummies equal to one for firms in the third and fourth quartile of the carbon intensity distribution. A firm's emission intensity is defined as the emission-weighted average emission intensity of its installations over the period 2013–2016 (see Section 4.1 Difference-in-differences setup in the paper for more details) All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.5: Heterogeneity by NACE 2-digit industry

Dependent Sample Specification	Emissions			
	Basic metals (24) (1)	Non-metallic minerals (23) (2)	Chemicals (20) (3)	All other (4)
$\text{Post}_t \times \text{Exp}_f$	-0.096*** (0.030)	-0.058*** (0.021)	-0.049 (0.037)	-0.028 (0.029)
Firm FE	Yes	Yes	Yes	Yes
Year-country FE	Yes	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes	Yes
Observations	2,263	9,676	2,252	6,151
Firms	206	881	205	561
Pseudo R^2	0.998	0.992	0.989	0.995

Notes: The sample consists of firm-year level observations for manufacturing firms for the period 2013–2023. The dependent variable is the firm’s emissions (in ton CO₂ equivalents). This specification examines heterogeneity in emission responses across major manufacturing subsectors. Post_t is a dummy equal to one if $t \geq 2017$, and zero otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. Samples are defined by NACE 2-digit industry codes: Basic metals (NACE 24) includes steel production; Non-metallic minerals (NACE 23) includes cement production; Chemicals (NACE 20) includes chemical manufacturing; All other includes all remaining 2-digit manufacturing industries. A firm’s exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications include firm fixed effects, year-country fixed effects, and year-industry fixed effects. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.6: Testing for Differential Treatment Effects: Orbis vs. Non-Orbis Firms

Dependent Sample Specification	Emissions		
	Full (1)	Power (2)	Manufacturing (3)
$\text{Post}_t \times \text{Exp}_f$	-0.122*** (0.037)	-0.344*** (0.080)	-0.057*** (0.017)
NotOrbis_{ft}	0.013 (0.082)	0.064 (0.203)	-0.003 (0.032)
$\text{Exp}_f \times \text{NotOrbis}_{ft}$	-0.051 (0.097)	-0.181 (0.217)	0.052 (0.049)
$\text{Post}_t \times \text{NotOrbis}_{ft}$	-0.033 (0.103)	-0.017 (0.223)	-0.033 (0.027)
$\text{Post}_t \times \text{Exp}_f \times \text{NotOrbis}_{ft}$	-0.042 (0.109)	-0.122 (0.231)	0.028 (0.034)
Firm FE	Yes	Yes	Yes
Year-country FE	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes
Fuel price controls	Yes	Yes	
Observations	24,202	3,738	20,430
Firms	2,207	343	1,861
Pseudo R^2	0.986	0.975	0.994

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm’s emissions (in ton CO₂ equivalents). This specification tests whether the treatment effect differs between firms with and without validated Orbis balance sheet data. Post_t is a dummy equal to one if $t \geq 2017$, and zero otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. NotOrbis_{ft} is a dummy equal to one for firm-year observations where emissions data are available but cannot be matched to validated Orbis operating revenue data, and zero otherwise. The reference group ($\text{NotOrbis}_{ft} = 0$) consists of firm-years with Orbis-matched data. A firm’s exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.7: Separate post dummies for 2017–2019 and 2020–2023

Dependent Sample Specification	Full (1)	Emissions Power (2)	Manufacturing (3)	Emissions / Operating Revenue Full (4)	Power (5)	Manufacturing (6)
$\text{Post}_t^{2017-2019} \times \text{Exp}_f$	-0.099*** (0.026)	-0.335*** (0.055)	-0.036*** (0.014)	-0.074*** (0.021)	-0.218*** (0.054)	-0.059*** (0.018)
$\text{Post}_t^{2020-2023} \times \text{Exp}_f$	-0.193*** (0.034)	-0.849*** (0.117)	-0.064*** (0.020)	-0.110*** (0.033)	-0.480*** (0.103)	-0.053* (0.032)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Fuel price controls	Yes	Yes		Yes	Yes	
Observations	24,202	3,738	20,430	15,848	2,443	13,338
Firms	2,207	343	1,861	1,619	245	1,368
Pseudo R^2	0.986	0.976	0.994	0.618	0.621	0.594

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the firm's emissions (in ton CO₂ equivalents) in columns 1–3, and firms' emission (in)efficiency, proxied by the ratio of total emissions (in tonne CO₂ equivalent) over operating revenue (in thousand euro) in columns 4 to 6. $\text{Post}_t^{2017-2019}$ is a dummy equal to one if $2017 \leq t \leq 2019$, and zero otherwise. $\text{Post}_t^{2020-2023}$ is a dummy equal to one if $2020 \leq t \leq 2023$, and zero otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm's exposure to the carbon price shock is defined at the firm level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation's verified emissions. All specifications are estimated with the Poisson Quasi-Maximum Likelihood estimator. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.8: Alternative estimators

Dependent Estimation	Emissions			
	Poisson	Poisson (weighted)	OLS	WLS
$\text{Post}_t \times \text{Exp}_f$	-0.136*** (0.027)	-0.134*** (0.028)	-0.058*** (0.021)	-0.054*** (0.020)
Firm FE	Yes	Yes	Yes	Yes
Year-country FE	Yes	Yes	Yes	Yes
Year-industry FE	Yes	Yes	Yes	Yes
Fuel price controls	Yes	Yes	Yes	Yes
Observations	23,171	23,171	23,171	23,171
Firms	2,207	2,207	2,207	2,207
Pseudo R^2	0.988	0.988		
R^2			0.952	0.952

Notes: The sample consists of firm-year level observations with non-zero emissions for the period 2013–2023. The first two columns report the unweighted Poisson QML estimates (column 1) and weighted Poisson QML estimates with weights given by a firm’s average emissions over the period 2013–2016. The second two columns report OLS (column 3) and WLS (column 4) estimates with the logarithm of emissions as a dependent variable. Post_t is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016, i.e., before the regulatory tightening. The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.9: Green deals with PQML

Dependent Sample Specification	Green Deals				
	(1)	Full (2)	(3)	Power (4)	Manufacturing (5)
$Post_t$	0.113 (0.159)				
Exp_f	-0.450 (0.321)				
$Post_t \times Exp_f$	0.609*** (0.221)	0.648*** (0.237)	1.144*** (0.341)	0.979 (0.961)	1.611*** (0.485)
Firm FE		Yes	Yes	Yes	Yes
Year FE		Yes			
Year-country FE			Yes	Yes	Yes
Year-industry FE			Yes	Yes	Yes
Observations	955	955	955	316	482
Firms	133	133	133	39	83
Pseudo R^2	0.014	0.288	0.358	0.360	0.370

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the number of green M&A deals that a firm announces in a given year. We use the pretrained large language model ClimateBERT (Webersinke et al. 2022) on the text incorporated as “deal rationale” in ZEPHYR to classify an M&A as “green” or not. $Post_t$ is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated via Poisson Quasi-Maximum Likelihood (Correia, Guimarães, and Tom Zylkin 2020). The sample covers 955 firm-year observations across 133 firms, which is smaller than the baseline OLS sample because the `ppmlhdfc` command identifies and drops “separated” observations to avoid issues related to the nonexistence of quasi-maximum likelihood estimates (Correia, Guimarães, and Thomas Zylkin 2025, Proposition 1); as a consequence, firms that do not engage in any deals are dropped from the estimation sample. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.10: Number of green deals (intersection of ClimateBERT and alternative definition)

Dependent Sample Specification	Green Deals (Intersection)					
	(1)	Full (2)	(3)	(4)	Power (5)	Manufacturing (6)
Post_t	0.010*** (0.003)					
Exp_f	0.001 (0.003)					
$\text{Post}_t \times \text{Exp}_f$	0.011* (0.006)	0.011* (0.006)	0.010 (0.006)	0.012 (0.007)	-0.010 (0.045)	0.013 (0.008)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	20,269	20,269	20,269	20,269	2,819	17,405
Firms	1,847	1,847	1,847	1,847	257	1,586
R^2	0.002	0.410	0.412	0.440	0.493	0.430

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the number of green M&A deals that a firm announces in a given year. We combine our rule-based labels with the original ClimateBERT classification to classify an M&A as “green” or not. Post_t is a dummy equal to one if $t \geq 2017$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated via OLS. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table OA.11: Number of green deals - subperiods (intersection of ClimateBERT and alternative definition)

Dependent Sample Specification	Green Deals (Intersection)					
	(1)	Full (2)	(3)	(4)	Power (5)	Manufacturing (6)
$\text{Post}_t^{2017-2019}$	0.007* (0.004)					
$\text{Post}_t^{2020-2023}$	0.011*** (0.003)					
Exp_f	0.001 (0.003)					
$\text{Post}_t^{2017-2019} \times \text{Exp}_f$	0.000 (0.006)	0.000 (0.006)	0.002 (0.006)	0.004 (0.007)	-0.024 (0.049)	0.007 (0.007)
$\text{Post}_t^{2020-2023} \times \text{Exp}_f$	0.019** (0.008)	0.019** (0.008)	0.016** (0.008)	0.019** (0.009)	0.091 (0.058)	0.017* (0.010)
Firm FE		Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes			
Year-country FE				Yes	Yes	Yes
Year-industry FE				Yes	Yes	Yes
Fuel price controls			Yes	Yes	Yes	
Observations	20,269	20,269	20,269	20,269	2,819	17,405
Firms	1,847	1,847	1,847	1,847	257	1,586
R^2	0.003	0.410	0.413	0.440	0.495	0.430

Notes: The sample consists of firm-year level observations for the period 2013–2023. The dependent variable is the number of green M&A deals that a firm announces in a given year. We combine our rule-based labels with the original ClimateBERT classification to classify an M&A as “green” or not. $\text{Post}_t^{2017-2019}$ is a dummy equal to one if $2017 \leq t \leq 2019$, and 0 otherwise. $\text{Post}_t^{2020-2023}$ is a dummy equal to one if $t \geq 2020$, and 0 otherwise. Exp_f is a dummy equal to one for all firms with carbon price exposure above the median in their NACE 4-digit industry. A firm’s exposure to the carbon price shock is defined at the firm-level as the average emission intensity of its regulated installations over the period 2013–2016 (i.e., before the regulatory tightening). The emission intensity of a power plant is inferred from its fuel and technology as provided by the Global Energy Monitor. The emission intensity of manufacturing installations is defined as the ratio of its historical emissions level divided by its historical activity level as inferred from the harmonized free allocation methodology. When a firm owns multiple regulated installations, we compute a weighted average emission intensity using each installation’s verified emissions. All specifications are estimated via OLS. Robust standard errors (clustered at the firm level) are in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Acknowledgements

We are grateful to Pat Akey, Geraldo Cerqueiro, Ralph De Haas, Manthos Delis, Ivan Ivanov, Mirabelle Muuls, Geyao Zhang, and seminar participants at Banco de España, 1st Bocconi-CEPR workshop on Financial Intermediation and Corporate Finance, Sciences Po, HEC Paris, Ghent University, the European Central Bank, Adam Smith Business School, Goethe Universität Frankfurt, the sustainable banking conference Zürich, ASSA-IBEF, the conference on 'Social and Sustainable Finance: Bridging Methods, Policy and Practice' at Brunel Business School for useful comments. We also thank Eva Mulder for excellent research assistance. Klaas Mulier gratefully acknowledges the financial support from the National Bank of Belgium and the FWO for a research grant G039723N. An earlier version of the paper was prepared for the 2020 Colloquium "Climate Change: Economic Impact and Challenges for Central Banks and the Financial System" at the National Bank of Belgium, and circulated as "Going green by putting a price on pollution: Firm-level evidence from the EU".

The views expressed do not necessarily reflect those of the European Central Bank, the National Bank of Belgium, or the Eurosystem.

Olivier De Jonghe

Nationale Bank van België/Banque Nationale de Belgique, Brussels, Belgium; European Central Bank, Frankfurt am Main, Germany; Tilburg University, Tilburg, The Netherlands; email: olivier.dejonghe@nbb.be

Klaas Mulier (corresponding author)

KU Leuven, Leuven, Belgium; email: klaas.mulier@kuleuven.be

Glenn Schepens

European Central Bank, Frankfurt am Main, Germany; email: glenn.schepens@ecb.europa.eu

Leonard Stimpfle

Ghent university, Ghent, Belgium; email: leonard.stimpfle@UGent.be

© European Central Bank, 2026

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network](#) electronic library or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-7932-0

ISSN 1725-2806

doi:10.2866/8867369

QB-01-26-174-EN-N