



HAL
open science

Cutting Free Allocation Cuts Emissions: Evidence from Phase IV of the EU ETS (Working Paper)

Arthur Willemaers

► **To cite this version:**

Arthur Willemaers. Cutting Free Allocation Cuts Emissions: Evidence from Phase IV of the EU ETS (Working Paper). 2025. <hal-05400137>

HAL Id: hal-05400137

<https://hal.science/hal-05400137v1>

Preprint submitted on 5 Dec 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Cutting Free Allocation Cuts Emissions: Evidence from Phase IV of the EU ETS

Working Paper

Arthur Willemaers^{1,2}

¹ University of Pau and the Adour Region, France

² HEC Liège, University of Liège, Belgium

Abstract

Phase IV (2021 to 2030) of the EU ETS tightened the cap and reformed free allocation. In 2021, a narrower carbon leakage list and benchmark updates produced a one-off drop in realized free allocation. Using 2015 to 2023 data on Phase III and Phase IV incumbents from the EUTL, we construct a time-invariant, installation-level dose from the 2020 to 2021 change and estimate a continuous difference in differences with installation and year fixed effects. The event study shows stable pre trends, and allocation levels are essentially flat after 2021.

Cutting free allocation reduces emissions. In our preferred specification, the slope after 2021 is -0.359 , implying about -3.5% lower emissions for a $+10$ percentage-point larger loss of free allocation. From the perspective of the independence property, this dose response pattern indicates that, in Phase IV manufacturing with activity-linked free allocation, changes in the allocation schedule are not neutral for verified emissions. Among exposed firms with positive pre-2020 banks, we find no detectable attenuation. Aggregating to the BvD-entity perimeter, the emissions response replicates in direction and timing. In standard Orbis accounting aggregates over 2021–2023 we do not detect short-run movements within the ETS-linked perimeter. Non-ETS plants of the same entity are unobserved, so we neither establish nor rule out carbon leakage. The evidence is consistent with stronger effective carbon pricing at the installation margin under activity-linked allocation in Phase IV and provides quantitative input on the environmental effect of phasing down free allocation and the transition to CBAM.

JEL Codes : Q58, Q52, D22, H23.

Key Words : EU ETS, Cap-and-Trade, Free Allocation, Carbon Leakage, Firm Behavior, Environmental Policy

1 Introduction

Phase IV (2021–2030) of the EU Emissions Trading System (EU ETS) tightened the cap and reformed free allocation. A pivotal institutional change in 2021 was the revision of the carbon-leakage (CL) list and benchmark updates, which sharply reduced free allocation for many manufacturing incumbents. In textbook cap-and-trade with lump-sum allocation, the Coase/Montgomery independence property implies that emissions are invariant to the initial allocation: only the cap should matter. In the EU ETS manufacturing setting, however, free allocation is activity-linked (output-based allocation with dynamic updating), so changes in per-unit free allocation can affect marginal compliance costs. The Phase III→IV transition therefore offers a transparent design of whether cutting free allocation induces measurable abatement in practice.

Three features make this moment informative. First, Phase IV is the first period in which the industrial ETS is tightly binding under sustained high EUA prices, so abatement incentives from allocation changes are economically meaningful. Second, the 2021 CL-list/benchmark reform generated large, plausibly exogenous cross-sectional cuts in free allocation across otherwise comparable installations. Third, administrative allocation is empirically stable after 2021, so the reform delivers a one-off “dose” realized at implementation rather than a gradual drift. Together these facts allow a transparent question: *does less free allocation reduce verified emissions, by how much, and for whom?*

This setting also sharpens the link to the independence property. Existing evidence for the EU ETS points to approximate neutrality for large emitters and departures for some smaller plants that are consistent with trading frictions and behavioral effects. In contrast, Phase IV manufacturing combines a tightly binding cap, sustained high allowance prices, and activity-linked allocation rules that tie free allocation directly to output and carbon-leakage status. The 2021 reform therefore shifts the effective carbon price faced by incumbents through an institutional channel, not only through frictions. Measuring how emissions respond to this one-off change provides a direct reduced-form test of whether independence holds under the Phase IV design.

We answer this using an installation-level, time-invariant treatment intensity built directly from observed free allocation around the 2021 break, and a difference-in-differences design with a continuous treatment. The main specification compares post-2021 changes in log emissions across

installations with higher versus lower exposure to the reform, absorbing installation fixed effects and year fixed effects. An event-study documents pre-trends and implementation dynamics, and we restrict the estimation sample to Phase III CL incumbents to avoid contamination from new-entrant and derogation mechanics. As an intertemporal complement, we aggregate to the firm level to test whether pre-2020 allowance banks attenuated the immediate post-reform response, scaling banks by one baseline year of expected shortfall. We also link installations to Orbis BvD entities to build an ETS-linked firm-year panel, which we use to re-estimate the design at the firm perimeter and to examine standard accounting outcomes.

Two results emerge. First, cutting free allocation reduces emissions: in our preferred continuous DiD (installation FE and year FE, excluding anticipation years 2019-2020), the post-2021 slope is -0.359 (SE 0.070), which implies about -3.5% lower emissions for a $+10$ p.p. larger loss of free allocation. Second, among exposed firms with positive pre-2020 banks, we find no detectable attenuation when scaling banks by need: the continuous moderator is essentially zero and the high-low buffer groups have indistinguishable post-2021 exposure slopes (difference 0.016, $p > 0.1$). At the firm perimeter, the emissions response replicates in direction and timing. In Orbis accounting aggregates over 2021–2023 we do not detect short-run movements; these nulls apply only within the ETS-linked perimeter and do not speak to reallocation between ETS and non-ETS plants of the same entity.

In this paper, we (i) construct an installation-level, time-invariant dose from administrative free-allocation data and verify that allocation is essentially unchanged after 2021, so the dose is realized at implementation, (ii) provide a clean continuous DiD around a single policy break in manufacturing, where activity-linked allocation means independence need not hold, (iii) deliver EU-wide, cross-sector semi-elasticities of emissions with respect to cuts in free allocation, expressed in policy-relevant units, (iv) test whether pre-2020 banking muted the initial abatement and find a precise null, and (v) corroborate the emissions slope at the BvD-entity level while reporting null short-run changes in standard Orbis accounting variables within the ETS-linked perimeter. Non-ETS plants of the same entity are unobserved, so we neither establish nor rule out carbon leakage. Taken together, the results inform the ongoing phase-down of free allocation and the transition toward CBAM without overstating what the cap alone guarantees.

2 Institutional background: the EU ETS in Phases III-IV

The EU ETS is a cap-and-trade program covering power, heat, and energy-intensive industry across EU and associated countries. A declining EU-wide cap limits total verified emissions; regulated entities surrender one allowance per tonne; allowances are obtained by auctioning and, for selected sectors, by free allocation.

Phase III unified the market under an EU-wide cap with a linear reduction factor, made auctioning the default, and maintained free allocation for sectors on a carbon-leakage (CL) list (initial list 2013-2014; revised list 2015-2019, later extended to 2020). Product benchmarks were applied and gradually tightened, the Market Stability Reserve (MSR) was created to manage surplus allowances.¹

Phase IV tightens the cap and reforms free allocation.² Benchmarks are recalibrated and decline over time, the eligibility for full carbon-leakage protection is restricted to the revised list, and activity-level updating is explicitly codified. Even installations that remained on the list face a mechanical reduction in their free allocation because the new product benchmarks are lower and the cross-sectoral correction factor continues to bind. Non-listed sectors receive only a 30% share of benchmarked allocation in 2021, phasing down to zero by 2030. Together, these reforms generate a sharp cross-sectional drop in realized free allocation in 2021, with much larger losses for sectors that left the list.

3 Related Literature

This paper relates to the benchmark of allocation independence and to empirical work on the EU ETS. A first strand studies the independence property in cap and trade. In a standard setting with perfect trading and a binding cap, emissions are independent of the initial allocation, the Coasean independence property. Departures can arise through trading and transaction costs, collateral and credit constraints, risk and hedging motives, or institutional rules that tie allocation to activity through output based updating. Evidence from the EU ETS suggests neutrality for large emitters

¹Dynamic allocation rules operated already in Phase III for substantial activity changes; we use “dynamic allocation” generically to denote activity-linked updating.

²The Phase IV framework is set by Commission Delegated Regulation (EU) 2019/331 of 19 December 2018 and Commission Delegated Decision (EU) 2019/708 of 15 February 2019.

and departures for smaller ones that are consistent with frictions (Zaklan 2016; Zaklan 2023). Related work highlights disproportionately high trading costs for small firms (Jaraitė, Convery, and Di Maria 2010; Heindl 2012; Eden, De Oca, and Haug 2019; Löfgren et al. 2018). These mechanisms are likely to matter in Phase IV when EUA prices are high and when design features such as list tightening and activity updates become salient.

A second strand examines the effectiveness of the EU ETS. Early work on Phases I and II found mixed impacts on emissions (Ellerman and Buchner 2008; Abrell, Ndoye Faye, and Zachmann 2011; Jaraitė and Maria 2016). With richer micro data and stronger identification, more recent studies tend to find significant abatement. (Bordignon and Gamannossi degl’Innocenti 2023) estimate cumulative reductions of about 422 MtCO₂, roughly 22.5% of covered annual emissions, using event study and difference in differences across Phases II and III. (Dechezleprêtre, Nachtigall, and Venmans 2023) match EUTL and Orbis firm data in four countries and show emissions cuts without harm to firm performance. These results provide context for our setting, which focuses on a discrete institutional change within Phase IV rather than on system wide impacts across phases.

A related literature analyzes rules and behavioral incidence of free allocation. Reviews emphasize eligibility based on carbon and trade intensity, benchmark construction, operational effects, and adjustment mechanisms (Verde, Teixidó, et al. 2019). Several studies argue that leakage risk was overstated and that eligibility was too broad (Bruyn, Schep, and Cherif 2016; Ralf Martin et al. 2014a; Ralf Martin et al. 2014b; Sato et al. 2015). Performance based benchmarks introduced in Phase III reduced total free allocation and improved targeting (Sartor, Pallière, and Lecourt 2014), though sectoral heterogeneity can still yield mis targeting (Stenqvist and Åhman 2016). On behavioral incidence, (Ellerman and Reguant 2008) find no systematic production response to initial allocation. Phase III evidence points to neutrality for most emitters and to endowment effects for small plants (Zaklan 2016), and sectors deemed at risk abated less, consistent with undervaluing the opportunity cost of free allowances (De Vivo and Marin 2018). Allocation adjustment rules also shape behavior through closures and exit patterns (Verde, Graf, and Jong 2019), and activity thresholds spurred strategic output responses in cement (Neuhoff et al. 2014; Branger et al. 2015).

Our study sits at the intersection of these strands. We examine manufacturing in Phase IV,

where activity linked free allocation can affect marginal incentives, and we use an observed time invariant installation dose constructed from administrative free allocation data for 2020 and 2021 to study the emissions response under a binding regime.

Relative to the independence-property literature, we move from earlier phases and sectors to a setting in which the cap is tightly binding, allowance prices are high, and allocation rules are explicitly linked to activity and carbon-leakage eligibility. Zaklan (2016, 2023) document approximate neutrality for many large emitters and departures consistent with frictions in earlier phases. We complement this by exploiting the Phase III→IV reform as an institutional change in the allocation schedule for manufacturing incumbents and by estimating a continuous dose–response of emissions with respect to the realized cut in free allocation. In this environment, independence need not hold even in the absence of frictions.

Relative to work on the effectiveness of the EU ETS and on free allocation, we focus on a discrete design change within Phase IV rather than on phase-wide comparisons. We provide EU-wide, cross-sector estimates for manufacturing incumbents instead of selected countries, and we combine installation-level registry data with an Orbis-based firm panel to track both emissions and accounting aggregates at the ETS-linked perimeter. This lets us quantify the environmental effect of cutting free allocation while documenting, within scope, the absence of short-run movements in standard performance proxies.

4 Data

We assemble an installation-year panel from administrative EU ETS records for Phase III and Phase IV and enrich it with the Phase III and Phase IV carbon-leakage lists and with sectoral context from Eurostat. The combination of these sources allows us to define the outcome and the free-allocation dose at the installation level, to link installations to parent firms for clustering and banking, and to control for broad sectoral conditions where needed. We also link installations to Orbis Europe (`bvdid`) to build an ETS-only firm-year panel with standard accounting variables. The linkage lets us aggregate ETS installations to the `bvdid` perimeter and computing firm-level emissions and exposure using pre-period emissions-share weights and test accounting outcomes. These variables proxy complementary economic margins: activity and scale (employment, operating revenue, total

assets, fixed assets), operating performance (EBITDA, staff costs), and balance-sheet position (total and current liabilities). We use them later to characterize whether emission cuts coincide with output changes, cost pressure, or financing stress at the BvD-entity perimeter.

The registry data come from the European Union Transaction Log and include verified emissions, annual free allocation, installation identifiers, and compliance years. We merge the Phase III and Phase IV carbon-leakage lists to track eligibility for full versus partial protection and to date the 2021 revision. We add macro indicators from Eurostat to provide descriptive context and optional controls. Throughout, "free allocation" means incumbent free allocation under the benchmark, exposure and CSC rules. It excludes new-entrants allocations and Article 10c transitional allocations.³

The unit of observation is the installation by year. Each installation is linked to a parent-firm identifier so that we can account for multi-plant ownership. Standard errors are clustered at the firm level to allow arbitrary correlation across an owner's installations through time.

The outcome variable is the logarithm of verified emissions. Free allocation in levels is the annual number of freely allocated allowances recorded in the registry for each installation and year. The institutional decomposition into benchmarks, exposure factors, and the cross-sectoral correction factor is summarized in appendix 9.1 and is not required for the empirical construction used here.

Treatment intensity is defined as the one-off drop in free allocation at the transition from Phase III to Phase IV. For installation i the time-invariant dose is

$$\text{intensity}_i = 1 - \frac{FA_{i,2021}}{FA_{i,2020}} \in [0, 1].$$

Here, $FA_{i,2020}$ and $FA_{i,2021}$ refer to carbon leakage list incumbent free allocation (NER and Article 10c excluded). An increase of 0.1 indicates a ten percentage point larger loss of free allocation from 2020 to 2021. Because benchmark and correction factor movements are largely common within a year and activity updates after 2021 are small, cross sectional variation in this dose is primarily driven by the change in carbon leakage status in 2021. In the estimation sample that includes Phase III carbon leakage incumbents, installations that remain protected have doses close to zero, whereas

³Incumbent FA excludes new-entrants reserve (NER) allocations and Article 10c transitional allocations. See Appendix9.1 for details .

those that lose protection have doses around 0.7, indicating a shift from almost full protection to roughly thirty percent. Free allocation remains empirically stable after 2021, so the dose is realized at implementation. Appendix Table 3 reports percentage changes in free allocation for 2022 and 2023 relative to 2021, with medians at zero and narrow interquartile ranges.

The sample spans 2015 to 2023 so that the pre- and post-reform periods align with the relevant institutional regimes. The Phase III carbon-leakage list was updated in 2015 and applied through 2020 and 2021 marks the start of Phase IV with the list and benchmark revision. Years 2013 and 2014 are not used in order to maintain a consistent pre-policy environment. The year 2015 is a transitional year that reflects first-year accommodation to the updated Phase III scope and it does not anchor identification. It is sometimes excluded in specifications that emphasize a stable pre-period.

We restrict the estimation sample to installations that were on the Phase III carbon leakage list. This keeps the pre period under a single protection regime and excludes facilities that could receive free allowances through new entrant provisions or special derogations. Those allocation streams arise for reasons unrelated to the 2021 list revision and would add noise to the dose based on the 2020 to 2021 change. The restriction ensures that comparisons are between like installations within sectors and within years. Following common practice at this level of aggregation, and as in (Zaklan 2023), we implement the analysis at NACE level 1 and retain only sectors with adequate support, defined as at least fifty installations and non zero dispersion in the dose so that comparisons occur among facilities operating in the same broad market environment within a given year. Observations with undefined or inconsistent allocation data are dropped. Appendix Table 4 lists the filters that generate the final sample.

For the intertemporal analysis, we construct a pre 2020 allowance stock at the firm level by cumulating annual surpluses through 2019, truncating at zero, and scaling by one baseline year of the expected shortfall implied by the portfolio weighted dose. This procedure yields a moderator expressed in years of coverage, and we use it later to test whether banking muted the immediate post reform response. Both the intensity and the portfolio weights used later are computed from Phase III carbon leakage incumbent installations and their incumbent *FA* series.

4.1 Descriptive statistics

The panel covers 2,253 installations operated by 1,617 parent firms over 2015 to 2023. Roughly one fifth of parents are multi-plant yet they own about two fifths of installations, so we cluster at the parent firm. Table 1 summarises the panel structure we use throughout.

Figure 1 shows the cross-sectional distribution of the dose. A large mass sits near 0.2 for installations that remained on the protection list and a second mass reflects the large losses among leavers. The histogram is reported in Appendix Figure 3. The group split is documented in Appendix Table 8. Median intensity is 0.209 for stayers and 0.738 for leavers, as reported in Appendix Table 9.

Emissions decline from 2018 to 2023, with a temporary rise in 2021⁴, while free allocation drops in 2021 and is essentially stable thereafter. Appendix Table 5 reports installation-level means, medians, and dispersion for both series. System-wide free allocation relative to emissions falls from about 1.10 in 2020 to 0.80 in 2021 and rebounds to 0.94 by 2023. Appendix Table 6 provides the ratio of totals and the distribution of installation-level ratios. For completeness, Appendix Table 7 reports detailed summary statistics for the dose in the analysis sample. Stability of free allocation after 2021 is shown in Appendix Table 3. The step-by-step sample construction and filters appear in Appendix Table 4.

These facts confirm that the reform created wide cross-sectional exposure while the post-2021 allocation path is stable, which supports the continuous treatment design developed next.

Table 1: Panel structure summary

Installations (unique <code>inst_num</code>)	2,253
Parent firms (clusters, <code>firm_num</code>)	1,614
Years covered	2015–2023 (9 years)
Share of multi-plant firms	19.7% (mean installations per firm: 1.40)
Share of installations in multi-plant firms	42.5%

Notes: Multi-plant firm = parent with ≥ 2 installations in the panel. Roughly 80% of parents are single-plant, while the remaining $\sim 20\%$ (multi-plant) own about 43% of installations. This nesting makes within-firm correlation likely, so we cluster standard errors at the parent-firm level (1 614 clusters).

⁴COVID-19 rebound effect

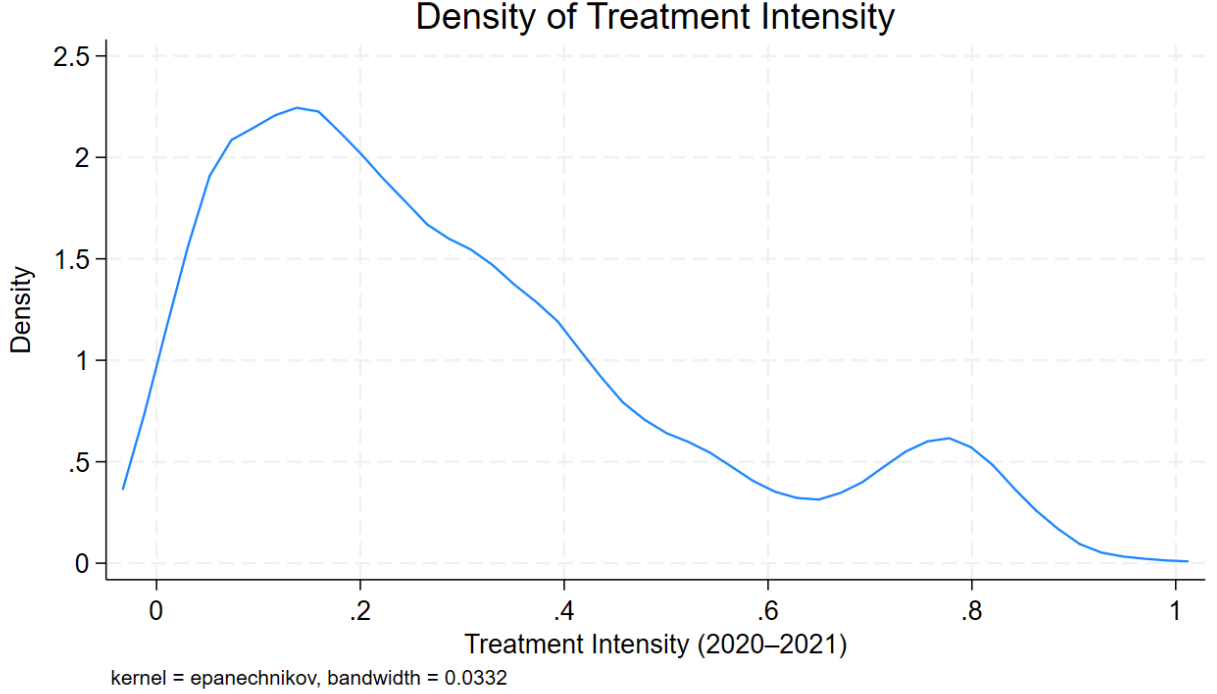


Figure 1: Kernel density of treatment intensity $1 - FA_{2021}/FA_{2020}$.

5 Empirical Strategy

Identification comes from the post versus pre contrast within the same installation, where $Y_{it} = \ln E_{it}$ is log verified emissions. The event study in Figure 2 shows a flat pre period by intensity, and the pre-trend slope test in Appendix Table 19 confirms that slopes do not vary with the dose.

Free allocation (FA) follows

$$FA_{it} = \theta_t \cdot HAL_{i,t},$$

where HAL is the historical activity level⁵ and $\theta_t = BM_t \cdot CLEF_t \cdot CSCF_t$. Here BM_t are product benchmarks, $CLEF_t$ is the carbon-leakage exposure factor equal to one on the Phase IV list and 0.3 otherwise with a scheduled phase-down, and $CSCF_t$ is the cross-sectoral correction factor. See Appendix 9.1 for details.

We define a time-invariant installation dose as intensity $_i \equiv 1 - FA_{i,2021}/FA_{i,2020}$. An increase of

⁵HAL is defined in Art. 15 of the FAR. For 2021–2025 it is the arithmetic mean of annual activity over the 2014–2018 baseline (Commission Delegated Regulation (EU) 2019/331, Art. 15(3)–(6)).

+0.10 corresponds to a ten-percentage-point larger loss of free allocation at the 2020–2021 break. Benchmark updates and the $CSCF_t$ move within years and are absorbed by year fixed effects, so cross-sectional variation in intensity $_i$ comes from the 2021 carbon-leakage list revision. Allocation is essentially flat after 2021, so the dose is realized at implementation (Appendix Table 3). We estimate dynamics with an event study and the average effect with a continuous difference-in-differences.⁶

5.1 Event study

With 2018 omitted as the pre-announcement baseline, we estimate

$$Y_{it} = \alpha_i + \sum_y \gamma_y \mathbf{1}\{t = y\} + \sum_y \beta_y (\mathbf{1}\{t = y\} \times \text{intensity}_i) + u_{it}, \quad (1)$$

where α_i are installation fixed effects and γ_y are year indicators. Standard errors are clustered at the parent-firm level. Here β_y are the intensity-slopes by year relative to 2018. The event study checks flat pre-trend interactions in 2016 and 2017, assesses anticipation in 2019 and 2020, and contrasts the 2021–2023 average with the anticipation years.

5.2 Average post effect

Our headline specification is

$$Y_{it} = \alpha_i + \lambda_t + \beta(\text{Post}_t \times \text{intensity}_i) + \varepsilon_{it}, \quad (2)$$

where $\text{Post}_t \equiv \mathbf{1}\{t \geq 2021\}$, α_i are installation fixed effects, and λ_t are year fixed effects. Standard errors are clustered at the parent-firm level. Because the event study shows anticipation around the late-2018 policy announcement, we also report an implementation-only contrast that omits 2019–2020. A +0.10 increase in the dose implies a change in emissions of $100[\exp(0.10 \hat{\beta}) - 1]$ percent.

5.3 Identification and estimation

We restrict to Phase III carbon-leakage incumbents to remove new-entrant and derogation streams unrelated to the 2021 list revision. Parallel trends require that pre-policy slopes do not vary with the dose; the event study and a direct pre-trend slope check confirm this. We exclude 2015. The new Phase III carbon-leakage list for 2015–2019 entered into force on 1 January 2015, altering which

⁶Year fixed effects absorb common COVID dynamics. A spurious correlation would require installation-level shocks that co-vary with the dose. The flat pre period and the slope test in Appendix Table 19 do not support such a pattern.

sectors received full benchmarked free allocation; the list then remained in place through 2020. We therefore use 2016–2018 as the stable pre period.

Following (Zaklan 2023), we estimate at NACE Level 1 and retain only sectors with sufficient support, defined as at least fifty installations and nonzero dispersion in the treatment intensity, so comparisons occur within a comparable market environment in each year. Level 1 avoids small-cell problems that appear at finer NACE levels, where the 2021 carbon-leakage revision was applied at the Level 3 granularity and left most Level 3 cells internally homogeneous. At Level 2, the distribution of treatment intensity remains uneven and support is patchy. Level 1 therefore strikes a balance: it preserves within-sector cross-sectional dispersion in the 2020 to 2021 allocation cut while maintaining meaningful economic comparability across firms. Robustness to sector composition is confirmed by a leave-one-sector-out exercise in Appendix Table 12.

We re-estimate Eqs. (1) and (2) at the BvD-entity level by replacing installations i with entities f . Outcomes Y_{ft} are either ln of ETS emissions aggregated across the entity’s linked installations or Orbis accounting variables. Variables that are strictly positive enter in logs, while variables that can take zero or negative values enter in inverse hyperbolic sine. The dose remains intensity with subscript f , constructed as an emissions-share aggregation of installation doses within each entity using pre-period average verified emissions over 2015-2019. The mapping uses registry installation links to Orbis BvD identifiers.

5.4 Estimation window

The main estimation window covers 2016–2018 and 2021–2023, excluding 2019 and 2020 as anticipation years. The event study is shown over the full 2016–2023 range for continuity, and the 2015 extension appears in Appendix Figure 5 to document the one-year adjustment after the Phase III list revision. This hierarchy makes explicit that our difference-in-differences design identifies the post-reform change from 2021 onward, net of any pre-adjustment or anticipation behavior around the Phase IV announcement.

6 Results

6.1 Event study

Figure 2 plots the coefficients from Eq. (1) with 2018 as the omitted year and reports their numerical values in Table 10. Pre trends are flat in 2016 and 2017 ($p = 0.472$ for equality). From 2019 onward the slope turns negative (2019: -0.073 , $p = 0.014$ and 2020: -0.259 , $p < 0.001$), which is consistent with anticipation around the Phase IV regime. Post implementation effects are larger and persistent: -0.251 in 2021, -0.392 in 2022, and -0.393 in 2023 (all $p < 0.001$). A formal contrast shows the post-2021 average is more negative than the anticipation years by -0.179 log points ($p = 0.002$). Extending the window to 2015 produces a visible one-year deviation and breaks the joint pre-trend test, which supports treating 2015 as transitional and excluding it from preferred estimates (see Appendix Figure 5). For magnitude, a ten-percentage-point larger loss of free allocation implies about 3.8% lower emissions in 2022.⁷

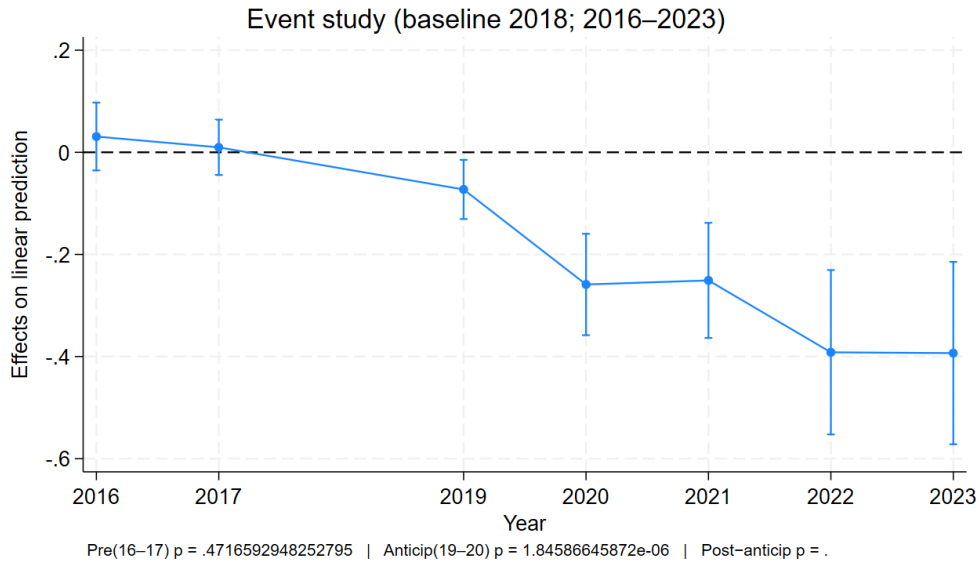


Figure 2: Event study (baseline 2018, 2016–2023). Coefficients are interactions of year dummies with treatment intensity, installation FE, SEs clustered by parent firm.

⁷Percent changes from log coefficients use $100 \times (\exp(0.10 \hat{\beta}_y) - 1)$ for a +0.10 change in intensity.

6.2 Average post effect

Table 2 reports estimates of Eq. (2). Column (1), our baseline over 2016–2023, estimates a $\text{post} \times \text{intensity}$ coefficient of -0.287 (SE 0.063). The implementation-only window that drops 2019–2020 (col. 2) yields -0.359 (SE 0.071). Replacing year fixed effects with $\text{sector} \times \text{year}$ fixed effects (col. 3) gives -0.432 (SE 0.094). Interpreted per +10 percentage points of intensity, these imply about 2.8%, 3.5%, and 4.3% lower emissions after 2021, respectively.⁸ For completeness, Appendix Table 11 shows the full-window estimate including 2015: -0.316 (SE 0.064). Because 2015 is a transitional year following the Phase III CL-list update, we exclude it from preferred specifications. These semi-elasticities summarize the reduced-form response and do not distinguish output from intensity changes. Adding $\text{country} \times \text{year}$ fixed effects in column (4) yields a virtually identical slope of -0.359 (SE 0.070). This implies about 3.6% lower emissions per +10 p.p. of intensity, which indicates the result is not driven by national macro shocks that vary over time.

Table 14 augments the continuous DiD with macro controls that are plausibly relevant, namely GDP, investment, goods trade, and non-household electricity prices, and in column (6) the EUA spot price. Across all three specifications the post reform dose response remains negative, stable, and precisely estimated. The coefficients on $\text{post} \times \text{intensity}$ are -0.374 (SE 0.086), -0.448 (SE 0.114), and -0.371 (SE 0.086) in columns (4) to (6). These imply roughly 3.7%, 4.4%, and 3.6% lower emissions, respectively, for a +10 p.p. increase in exposure.⁹ The result is robust when replacing common year fixed effects with $\text{sector} \times \text{year}$ fixed effects in column (5), and when adding the EUA price in column (6). Introducing the EUA series requires dropping year fixed effects to avoid collinearity, yet the $\text{post} \times \text{intensity}$ slope is virtually unchanged. Among controls, the EUA coefficient is negative and precisely estimated, the electricity price coefficient is small and imprecise once rich fixed effects are included, and trade covariates do not overturn the treatment effect. Investment is negligible and statistically indistinguishable from zero when time fixed effects or $\text{sector} \times \text{year}$ fixed effects are present, that is columns (4) and (5). It becomes negative and significant only in the specification without time fixed effects that includes the EUA price, that is column (6). This pattern is consistent with a possible capital renewal channel, but it also cautions that broader macro shocks may load onto this coefficient. Overall, the controlled specifications corroborate the main finding of a persistent

⁸Percent changes from log coefficients use $100 \times (\exp(0.10 \hat{\beta}) - 1)$ for a +0.10 change in intensity.

⁹Percent changes from log coefficients use $100 \times (\exp(0.10 \hat{\beta}) - 1)$.

post 2021 contraction in emissions that is proportional to the loss of free allocation.

Table 2: Average post effect (continuous DiD): Log emissions vs. treatment intensity

	(1)	(2)	(3)	(4)
Post × Intensity	-0.287***	-0.359***	-0.432***	-0.359***
	(0.063)	(0.071)	(0.094)	(0.070)
Installation FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes		
Sector×Year FE			Yes	
Country×Year FE				Yes
Observations	18,024	13,518	13,518	13,518
Clusters (firm)	1,614	1,614	1,614	1,614
Within R^2	0.0072	0.0105	0.0106	0.0105

Notes: Dependent variable is $\ln(\text{emissions})$. **(1)** Baseline TWFE with installation and year fixed effects on 2016–2023. **(2)** Implementation-only window (drops 2019–2020), same FE as (1). **(3)** Implementation-only; replaces year FE with sector (NACE L1)×year FE. **(4)** Implementation-only; replaces year FE with country×year FE. Standard errors in parentheses, clustered at the parent-firm level. The main effects of intensity (time-invariant) and post (year-invariant) are absorbed by fixed effects and omitted due to collinearity. “Post × Intensity” corresponds to $\text{post}_t \times \text{intensity}_i$. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Robustness and heterogeneity

The dose-response is monotone across the treatment distribution. A non-parametric quintile specification shows increasingly negative post effects from the lower to the upper bins (Appendix Table 15), and at the top quintile the post coefficient is around -0.25 relative to the lowest bin. A simple continuous model evaluated at bin means tracks these magnitudes closely (Appendix Table 16), which supports a linear-in-intensity summary. An alternative identification using the De Chaisemartin–D’Haultfœuille estimator yields a negative post effect and passes the quasi-untreated-group test (Appendix Tables 17–18). Finally, 2016–2018 pre-trend slopes are essentially orthogonal to the dose (Appendix Table 19), reinforcing the parallel-trends assumption.

A final timing check shows that the post 2021 slope is stable across different post windows.

We re-estimate the preferred specification keeping only 2021, then 2021 to 2022, then 2021 to 2023. The estimated slopes range from about -0.26 to -0.36 with overlapping confidence intervals. This stability indicates that the effect is not driven by a particular post year and supports the interpretation that emissions decline more where the loss in free allocation is larger (Appendix Table 13).

We also verify that no single sector drives the result. A leave one sector out exercise at NACE Level 1 re runs the preferred specification while excluding each sector in turn. The post times intensity slope remains negative and close to the baseline throughout. The range is -0.404 to -0.309 across exclusions. Appendix Table 12 reports the sector by sector values.

6.4 Banking and attenuation

We test whether pre-2020 banking muted the post-2021 exposure effect using two firm-year specifications. Specification A builds the panel from Phase III carbon-leakage (CL) installations only. Specification B aggregates emissions across all installations owned by the same firms, while keeping the same firm-level exposure. Both include firm and year fixed effects, and standard errors are clustered by firm.

Our estimating equation is

$$\ln E_{ft} = \alpha_f + \lambda_t + \beta_1(\text{Post}_t \times \text{TreatExp}_f) + \beta_2(\text{Post}_t \times \text{BufferNeed}_f) + \varepsilon_{ft}. \quad (3)$$

The exposure TreatExp_f is the firm’s portfolio-weighted installation dose computed from CL incumbents. The moderator

$$\text{BufferNeed}_f \equiv \frac{\text{bank}_{f,2019}^{\text{capped}}}{\text{TreatExp}_f \cdot E_{\text{base},f}}$$

scales the firm’s end-2019 bank by one baseline year of the expected shortfall. A value of one means the firm could fully cover one baseline year of the loss. We cap the bank at 2019 so the moderator is strictly pre-policy and uncontaminated by 2020. Portfolio weights use average emissions over 2015–2019 so that both TreatExp_f and $E_{\text{base},f}$ reflect the same pre period. Zero-bank firms are retained and assigned $\text{BufferNeed}_f = 0$ whenever the denominator is well defined. A median-split variant of BufferNeed_f is reported in Appendix Eq. (5). Construction details appear in Appendix 9.5.

Across specifications we find a clear post-2021 exposure effect with no evidence that pre-2020 banking attenuates it. In Specification A (CL-only panel) the post \times exposure slope is -0.293 with

standard error 0.070, and the bank-ratio moderator is near zero and imprecise (4.3×10^{-6} with standard error 8.5×10^{-6}), see Table 20. A median split of the buffer yields post slopes of -0.282 for low buffers and -0.335 for high buffers, with a small and insignificant difference of -0.053 (standard error 0.143). In Specification B (all-installations totals, excluding 2015) the post \times exposure slope remains negative and precise at -0.264 with standard error 0.060, and the bank-ratio moderator is again near zero and imprecise (3.16×10^{-6} with standard error 8.34×10^{-6}), see Table 21. The median-split post slopes are -0.198 and -0.340 , with an insignificant difference of -0.142 (standard error 0.124). For a $+0.10$ increase in exposure these estimates imply about 2.9% lower emissions in the CL-only panel and 2.6% lower emissions in the all-installations panel.¹⁰ We present both tables for clarity. Table 20 targets the policy-relevant CL segment used to construct exposure. Table 21 verifies that the same firm-level exposure predicts total firm emissions across the broader asset base. The similar slopes and uniformly null banking interactions indicate no attenuation from pre-2020 stocks in either construction.

6.5 Firm-level aggregation and Orbis outcomes

We re-estimate Eqs. (1) and (2) at the BvD-entity level over the implementation-only window (2016–2018 vs. 2021–2023) using BvD-entity fixed effects and NACE Level 1 \times year fixed effects, with standard errors clustered by BvD entity. The resulting firm-level estimates confirm the installation-level results. The post \times intensity coefficient on \ln firm emissions is -0.273 (SE 0.101), which implies about 2.7% lower emissions for a +10 percentage-point increase in exposure.¹¹ Using Orbis accounting outcomes (employment, operating revenue, total assets, total fixed assets, EBITDA, staff costs, total liabilities, current liabilities), post \times intensity coefficients are small and not statistically significant. We report the emissions result and the full panel for accounting outcomes in Appendix Table 22. These firm-level results pertain to the ETS-linked perimeter constructed from Orbis. The accounting nulls should be read as no detectable change in these proxies over 2021–2023 within this linked sample, not as proof of the absence of any reallocation response.

¹⁰Percent change for an exposure change Δ uses $100 \times (\exp(\hat{\beta} \Delta) - 1)$.

¹¹Percent change uses $100 \times (\exp(0.10 \times \hat{\beta}) - 1)$.

7 Discussion and policy implications

We find a clear dose response within the EU ETS. Installations that lost more free allocation in 2021 reduced verified emissions more thereafter. In our preferred implementation only specification the post reform semi elasticity is about -0.359 per unit of exposure, which implies roughly -3.6% for a $+10$ percentage-point increase in the loss of free allocation. See Table 2.

Timing aligns with the reform. The event study shows flat interactions in the stable pre period, then more negative slopes from 2019 onward and persistent effects after 2021. See Figure 2. Free allocation is empirically stable after 2021, which supports interpreting the 2021 change as a one-off dose realized at implementation rather than a gradual drift. See Appendix Table 3.

Firm-level aggregation and Orbis outcomes refine the reading. Re-estimating the design at the BvD-entity level yields a negative $\text{post} \times \text{intensity}$ slope for firm emissions that mirrors the installation result. In the same framework, standard accounting outcomes from Orbis: employment, operating revenue, total assets, total fixed assets, EBITDA, staff costs, total liabilities, and current liabilities do not display detectable movements over 2021–2023. These nulls should be read narrowly. The firm panel sums only ETS-linked installations within each BvD-entity, and non-ETS plants owned by the same entity are not observed. As a result, we do not observe total entity emissions, only ETS-linked emissions, and any reallocation of activity between ETS and non-ETS plants inside the same entity is outside our measurement. The firm-level exercise corroborates the emissions effect within the ETS-linked perimeter but neither establishes nor rules out carbon leakage. See Appendix Table 22.

Intertemporal buffers do not overturn the effect. Among exposed firms with positive pre-2020 banks the moderator is near zero, and high versus low buffer groups have statistically indistinguishable post slopes. See Appendix Table 20. This is consistent with a binding effective constraint even for buffered firms. It does not identify channels. From the perspective of the independence property, these reduced-form patterns indicate that once free allocation is tightly linked to activity and list eligibility in a binding Phase IV regime, emissions respond to the allocation schedule. The cap still constrains total emissions, but the cross-sectional distribution of free allocation shapes where abatement takes place and by how much.

Policy implications are correspondingly focused. Within the EU ETS, phasing down free allocation is associated with measurable reductions in verified emissions at covered installations. Over 2021–2023 and within the linked ETS perimeter, we do not detect systematic movements in standard accounting aggregates. This tempers short run competitiveness concerns in these proxies. It does not allow statements about carbon leakage. With the current data we neither establish nor rule out leakage outside or within entities.

Overall, our contribution is to quantify the post reform dose response in verified emissions using an installation-level, time-invariant exposure. We show alignment in timing. We corroborate the pattern at the firm perimeter while clarifying the measurement scope. Robustness to anticipation trimming, richer time controls and alternative estimators is documented in Table 2, Figure 2 and Appendix Tables 15 and 3.

8 Conclusion

We study the Phase IV reform of free allocation in the EU ETS as a quasi experimental shift in carbon price exposure. Exploiting cross sectional variation in the one off drop from 2020 to 2021, we estimate an event study and a continuous difference in differences. Emissions fall more at installations that lose more free allocation. In our preferred design, a +10 percentage-point larger loss is associated with about 3.5% lower emissions after 2021. Dynamics are monotone and the pattern is robust to alternative time controls, non-parametric bins, an alternative estimator, and added controls.

These results show that tightening free allocation has environmental bite for incumbent Phase III sectors. When fewer free allowances are granted, verified emissions decline more where the cut is larger. We also find no detectable attenuation among exposed firms with positive pre-2020 banks. Re-estimating the design at the BvD-entity perimeter yields a negative post times intensity slope for entity emissions that mirrors the installation result. In the same framework, standard Orbis accounting outcomes do not display detectable movements over 2021 to 2023.

Our scope is intentionally narrow. At the installation level we measure reduced-form effects on verified emissions and cannot separate output from intensity. At the firm perimeter we sum ETS-linked installations within each BvD-entity and do not observe non-ETS plants owned by the

same entity. As a result we do not observe total entity emissions and we cannot track reallocation of activity between ETS and non-ETS plants. We therefore neither establish nor rule out carbon leakage. The estimates inform the environmental margin of the phase-down of free allocation and the transition toward the border adjustment. More broadly, the results show that in Phase IV manufacturing, with high allowance prices and activity-linked allocation, the initial allocation is no longer neutral for verified emissions, which is a concrete and policy-relevant departure from the benchmark independence property.

Future work should add coverage of total firm activity across ETS and non-ETS assets and link to trade indicators. Applying the same continuous treatment framework to those data would allow a clean split between output and intensity and a direct assessment of displacement across borders. Until such data are available, our contribution is to quantify the post-reform dose response in verified emissions using an installation-level, time-invariant exposure and to show that it is stable across identification choices and robustness checks.

References

- Abrell, Jan, Anta Ndoye Faye, and Georg Zachmann (2011). *Assessing the impact of the EU ETS using firm level data*. Tech. rep. Bruegel working paper.
- Bordignon, Massimo and Duccio Gamannossi degl’Innocenti (2023). “Third time’s a charm? Assessing the impact of the third phase of the EU ETS on CO2 emissions and performance”. In: *Sustainability* 15.8, p. 6394.
- Branger, Frédéric et al. (2015). “EU ETS, free allocations, and activity level thresholds: the devil lies in the details”. In: *Journal of the Association of Environmental and Resource Economists* 2.3, pp. 401–437.
- Bruyn, Sander Maurits, N Schep, and Sofia Cherif (2016). *Calculations of Additional Profits of Sectors and Firms from the EU ETS*. CE Delft.
- De Vivo, Nicola and Giovanni Marin (2018). “How neutral is the choice of the allocation mechanism in cap-and-trade schemes? Evidence from the EU-ETS”. In: *Argomenti* 9, pp. 21–44.
- Dechezleprêtre, Antoine, Daniel Nachtigall, and Frank Venmans (2023). “The joint impact of the European Union emissions trading system on carbon emissions and economic performance”. In: *Journal of Environmental Economics and Management* 118, p. 102758.

- Eden, Alexander, Mariza Montes De Oca, and Constanze Haug (2019). “Striving to Keep ETS Simple”. In: *International Carbon Action Partnership (ICAP)*.
- Ellerman, A Denny and Barbara K Buchner (2008). “Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005–06 emissions data”. In: *Environmental and Resource Economics* 41, pp. 267–287.
- Ellerman, A Denny and Mar Reguant (2008). “Grandfathering and the endowment effect An Assessment in the context of the Spanish National Allocation Plan”. In:
- Heindl, Peter (2012). “Transaction costs and tradable permits empirical evidence from the EU emissions trading scheme”. In: *ZEW Discussion Papers* 12.
- Jaraitė, Jūratė, Frank Convery, and Corrado Di Maria (2010). “Transaction costs for firms in the EU ETS: lessons from Ireland”. In: *Climate Policy* 10.2, pp. 190–215.
- Jaraitė, Jūratė and Corrado Di Maria (2016). “Did the EU ETS make a difference? An empirical assessment using Lithuanian firm-level data”. In: *The Energy Journal* 37.2, pp. 68–92.
- Löfgren, Åsa et al. (2018). “Distribution of emissions allowances and the use of auction revenues in the European Union emissions trading system”. In: *Review of Environmental Economics and Policy*.
- Martin, Ralf et al. (2014a). “Industry compensation under relocation risk: A firm-level analysis of the EU emissions trading scheme”. In: *American Economic Review* 104.8, pp. 2482–2508.
- (2014b). “On the empirical content of carbon leakage criteria in the EU Emissions Trading Scheme”. In: *Ecological Economics* 105, pp. 78–88.
- Neuhoff, Karsten et al. (2014). *Carbon control and competitiveness post 2020: The cement report*. Climate Strategies.
- Sartor, Oliver, Clement Pallière, and Stephen Lecourt (2014). “Benchmark-based allocations in EU ETS Phase 3: an early assessment”. In: *Climate Policy* 14.4, pp. 507–524.
- Sato, Misato et al. (2015). “Sectors under scrutiny: evaluation of indicators to assess the risk of carbon leakage in the UK and Germany”. In: *Environmental and Resource Economics* 60, pp. 99–124.
- Stenqvist, Christian and Max Åhman (2016). “Free allocation in the 3rd EU ETS period: Assessing two manufacturing sectors”. In: *Climate Policy* 16.2, pp. 125–144.
- Verde, Stefano F, Christoph Graf, and Thijs Jong (2019). “Installation entries and exits in the EU ETS: patterns and the delay effect of closure provisions”. In: *Energy economics* 78, pp. 508–524.

Verde, Stefano F, Jordi Teixidó, et al. (2019). “Free allocation rules in the EU emissions trading system: what does the empirical literature show?” In: *Climate policy* 19.4, pp. 439–452.

Zaklan, Aleksandar (2016). “Free allocation and the endowment effect in cap-and-trade systems: evidence from the European electricity sector”. In.

— (2023). “Coase and cap-and-trade: evidence on the independence property from the European carbon market”. In: *American Economic Journal: Economic Policy* 15.2, pp. 526–558.

9 Appendix

9.1 Allocation definitions and stability

Definition of free allocation

FA_{it} is the annual number of freely allocated allowances (EUAs, tonnes CO₂) received by installation i in year t as recorded in the EUTL. By institutional design,

$$FA_{it} = \theta_t \cdot HAL_{i,t},$$

$$\theta_t = BM_t \cdot CLEF_t \cdot CSCF_t,$$

where HAL equals the historical activity level¹², BM_t is the product benchmark, $CLEF_t$ is the carbon leakage exposure factor, and $CSCF_t \in (0, 1]$ is a uniform cross sectoral correction factor applied if the free allocation cap binds. Under Phase IV, $HAL_{i,t}$ equals the historical activity level unless the two year average activity deviates by more than fifteen percent from HAL , in which case the year t free allocation is adjusted by the exact percentage change (Implementing Regulation (EU) 2019/1842, Articles 4–5).

Dose stability

Free allocation remains essentially unchanged after 2021. Table 3 reports percent changes relative to 2021 at the installation year level.

¹²HAL is defined in Art. 15 of the FAR: for 2021 to 2025 as the arithmetic mean of annual activity over the 2014 to 2018 baseline (Commission Delegated Regulation (EU) 2019/331, Art. 15(3) to (6))

Table 3: Stability of free allocation relative to 2021 (installation-year)

Year	% change vs 2021			
	p25	p50	p75	Mean
2021	0.00	0.00	0.00	0.00
2022	-0.90	0.00	0.06	0.04
2023	-4.94	0.00	3.70	0.22

Notes: % change is $100 \times (FA_{it} - FA_{i,2021})/FA_{i,2021}$. Values for 2022–2023 are computed at the installation-year level. Medians are 0 and interquartile ranges are narrow, indicating that free allocation is essentially unchanged after the 2021 reform. Figures are not winsorized. (For reference, $N = 2,250$ per year in this tabulation.)

9.2 Sample construction and descriptives

Sample construction

We document the steps from raw registry records to the final analysis sample. Table 4 lists each filter and the remaining count.

Table 4: Sample construction and filters

Step	Filter
1	Keep years 2015–2023
2	Keep Phase III CL incumbents
3	Keep NACE level 1 sectors with support ($N > 50$ and non zero treatment intensity variation)
4	Drop undefined allocation data
5	Drop inconsistent intensity ($\notin [0, 1]$)

Descriptive statistics

Table 5: Emissions and free allocation by year (2015–2023)

Panel A: Emissions (verified CO ₂ , tonnes)						
Year	N	Mean	Median	SD	p25	p75
2015	2,253	140,617.0	31,765.0	465,506.2	13,928.0	84,764.0
2016	2,253	139,786.8	31,913.0	469,074.7	14,216.0	83,403.0
2017	2,253	140,330.6	32,512.0	441,841.0	14,388.0	85,175.0
2018	2,253	139,863.0	32,509.0	441,918.4	14,609.0	86,403.0
2019	2,253	135,068.8	31,276.0	420,966.9	13,945.0	83,498.0
2020	2,253	123,431.2	28,881.0	377,629.9	12,734.0	76,570.0
2021	2,253	131,014.0	30,781.0	417,115.3	13,525.0	79,783.0
2022	2,253	122,739.6	27,987.0	396,608.8	12,141.0	72,643.0
2023	2,253	111,818.5	24,997.0	375,247.0	9,824.0	62,306.0

Panel B: Free allocation (EUAs, tonnes CO ₂)						
Year	N	Mean	Median	SD	p25	p75
2015	2,252	147,710.1	34,447.5	583,169.1	16,350.0	92,679.5
2016	2,252	147,339.9	33,739.5	584,234.0	16,059.0	90,671.5
2017	2,252	143,570.5	33,092.5	572,816.2	15,745.5	88,797.5
2018	2,252	139,245.9	32,389.0	542,832.1	15,447.5	87,510.0
2019	2,252	138,309.9	31,607.5	539,736.8	15,159.5	85,398.5
2020	2,253	135,466.3	30,849.0	529,042.5	14,878.0	83,558.0
2021	2,253	104,346.1	20,927.0	444,931.0	8,609.0	63,942.0
2022	2,253	104,836.1	20,653.0	459,075.0	8,481.0	62,255.0
2023	2,253	105,484.8	20,506.0	468,503.9	8,515.0	61,288.0

Notes: Percentiles are across installations. Emissions are verified CO₂ (EUTL). Free allocation is the number of freely allocated EUAs. Values are in tonnes of CO₂. 2020 reflects COVID-related disruptions; 2021 marks the start of Phase IV allocation rules.

Table 6: Free allocation relative to emissions (2018–2023)

Panel A: Ratio of totals ($\sum_i FA_i / \sum_i Emissions_i$)

Year	$\frac{\sum FA}{\sum EM}$
2018	0.995
2019	1.024
2020	1.098
2021	0.796
2022	0.854
2023	0.943

Panel B: Installation-level ratio FA_i/EM_i (median and interquartile range)

Year	N	Median	p25	p75
2018	2,252	0.946	0.768	1.223
2019	2,252	0.964	0.776	1.255
2020	2,253	1.028	0.803	1.365
2021	2,253	0.702	0.523	0.862
2022	2,253	0.754	0.551	0.968
2023	2,253	0.856	0.632	1.146

Notes: Panel A is the preferred systemwide measure (ratio of totals). Panel B summarizes the distribution of plant-level ratios via medians and IQR. Values computed from `allocatedfree_sum` and `verified_sum`. The 2020 values reflect COVID-19 disruptions; 2021 marks the Phase-IV allocation reform.

Table 7: Distribution of treatment intensity ($1 - FA_{2021}/FA_{2020}$)

	N	Mean	SD	Variance	Skew.	Kurt.	p_{25}	p_{50}	p_{75}
Treatment intensity	20,277	0.2970	0.2262	0.05119	0.9260	3.0313	0.1205	0.2417	0.4113

Notes: Stata `summarize`, `detail`. Statistics computed on the analysis sample (20,277 installation–year observations, 2015–2023). The treatment intensity is bounded in $[0, 1]$ by construction and equals one minus the ratio of free allocation in 2021 to that in 2020.

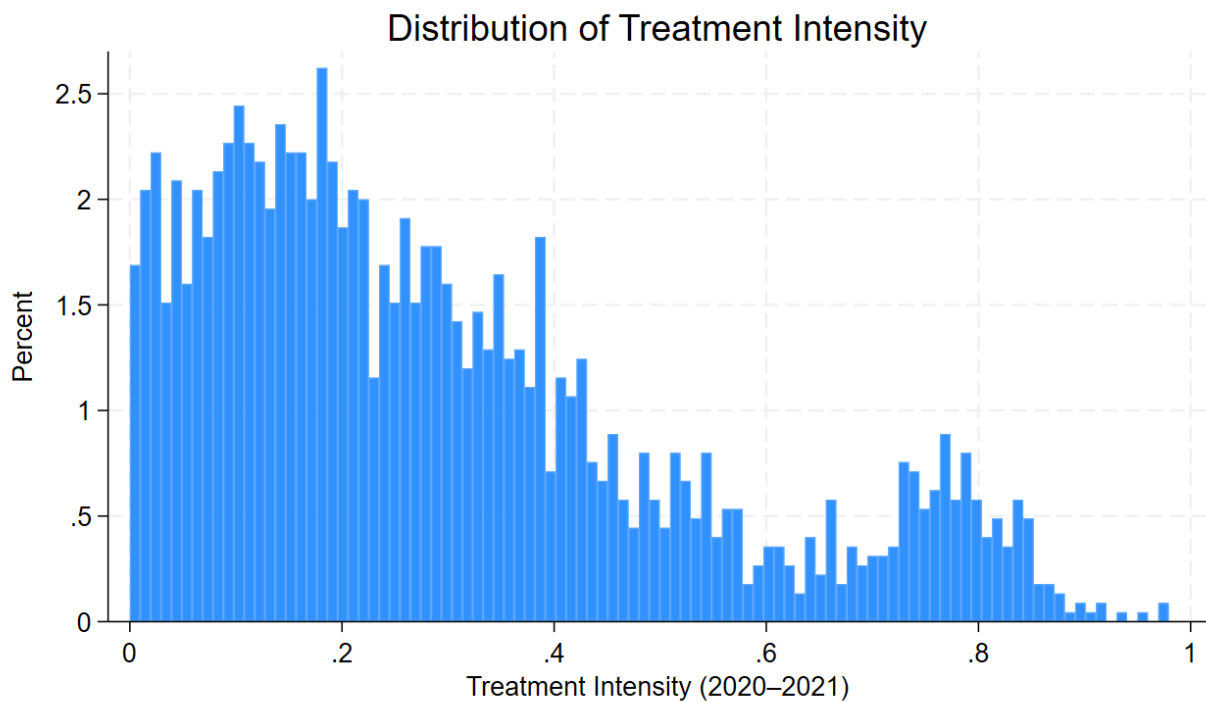


Figure 3: Histogram of treatment intensity $1 - FA_{2021}/FA_{2020}$.

Table 8: Transition groups from Phase III to Phase IV (installations)

Group	Installations	Share (%)
CL→CL	1,971	87.5
CL→non-CL	282	12.5
Total	2,253	100.0

Notes: Installation counts computed as $(\text{install.-years})/9$, since the balanced panel spans 9 years (2015–2023) after cleaning.

Table 9: Treatment intensity $(1 - \text{FA}_{2021}/\text{FA}_{2020})$ by transition group

Group	N	Median	p25	p75
CL→CL	1,971	0.209	0.108	0.349
CL→non-CL	282	0.738	0.568	0.792
All	2,253	0.242	0.121	0.411

Notes: Intensity is computed at the installation level as $1 - \text{FA}_{2021}/\text{FA}_{2020}$ using `allocatedfree_sum`; medians and quartiles are across installations within group. A +0.10 increase in intensity corresponds to a 10 p.p. larger loss of free allocation from 2020 to 2021.

9.3 Identification timeline

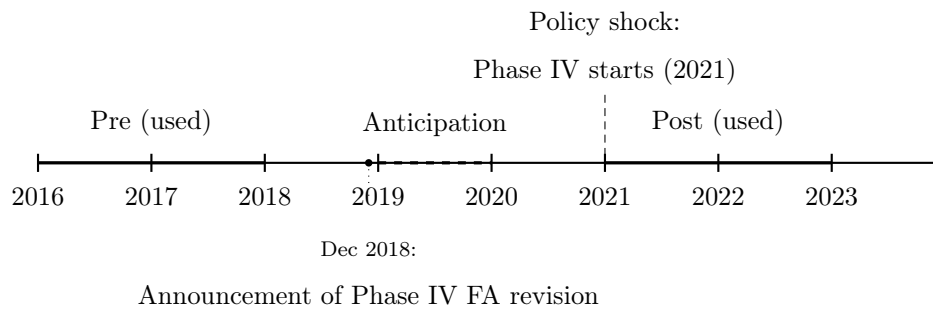


Figure 4: Identification timeline: pre (2016–2018) and post (2021–2023) periods. Anticipation years (2019–2020) are excluded in robustness checks. The dashed line marks the start of Phase IV (2021), and the dot at December 2018 indicates the announcement of the free allocation revision.

9.4 Additional results and diagnostics

Event-study coefficients

Table 10: Event-study of the Phase IV reform (baseline year 2018, firm-installation panel)

Year	$\hat{\beta}_t$	Std. Err.	<i>t</i> -stat	<i>p</i> -value
2016	0.0309	0.0339	0.91	0.363
2017	0.0097	0.0276	0.35	0.725
2019	-0.0728**	0.0295	-2.46	0.014
2020	-0.2589***	0.0507	-5.11	0.000
2021	-0.2509***	0.0575	-4.36	0.000
2022	-0.3917***	0.0821	-4.77	0.000
2023	-0.3933***	0.0911	-4.32	0.000
Pre (2016=2017) <i>p</i>		0.472		
Anticipation (2019,2020=0) <i>p</i>		0.000		
Post-Anticipation diff. <i>p</i>		0.002		
FE: installation		Yes		
Clusters (firm)		1,614		
Observations		18,024		
Within R^2		0.0943		

Notes: Baseline year is 2018. Coefficients are the marginal effects of TreatInt_f interacted with year dummies ($\text{Post}_t = \mathbb{1}\{t \geq 2021\}$). The regression absorbs installation fixed effects and clusters standard errors by firm. Pre-trend equality (2016 = 2017) cannot be rejected ($p = 0.47$), while anticipation effects (2019–2020 = 0) are jointly rejected ($p < 0.001$). The average post-2021 slope is significantly more negative than the anticipation period ($p = 0.002$). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Full window average post effect

Table 11: Average post effect (continuous DiD), full window 2015–2023

	(A1) 2015–2023
Post × Intensity	-0.316*** (0.064)
Installation FE	Yes
Year FE	Yes
Observations	20,277
Clusters (firm)	1,614
Within R^2	0.0082

Notes: Dependent variable is $\ln(\text{emissions})$. Specification matches Table 2 column (1) but includes 2015–2023. Standard errors in parentheses, clustered at the parent-firm level. Intensity and post main effects are absorbed by fixed effects (and omitted due to collinearity). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

9.4.1 Drop one NACE Level 1 sector

Table 12: Leave-one-sector-out robustness (preferred DiD specification)

Excluded NACE L1 sector	Post \times Intensity slope
All sectors (baseline)	-0.359***
Without sector 6	-0.389***
Without sector 8	-0.361***
Without sector 10	-0.352***
Without sector 11	-0.351***
Without sector 13	-0.361***
Without sector 17	-0.353***
Without sector 19	-0.335***
Without sector 20	-0.351***
Without sector 21	-0.362***
Without sector 22	-0.391***
Without sector 23	-0.404***
Without sector 24	-0.309***
Without sector 26	-0.366***
Without sector 27	-0.361***
Without sector 28	-0.341***
Without sector 29	-0.342***
Without sector 30	-0.375***
Range	[-0.404, -0.309]

Notes: Each row reports the coefficient on post \times intensity from the preferred specification (2016–2018; 2021–2023) when excluding that NACE Level 1 sector. All slopes remain negative and close to the baseline estimate (-0.359, SE 0.071). Standard errors are clustered at the firm level.

Activity Updating After 2021

Table 13: Post \times Dose slopes under trimmed post windows

	Dependent variable: ln(emissions)		
	(1)	(2)	(3)
	2021 only	2021 to 2022	2021 to 2023
Post \times Dose	-0.264*** (0.060)	-0.335*** (0.068)	-0.359*** (0.071)
Installation FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Clusters (firm)	1,614	1,614	1,614
Observations	9,012	11,265	13,518

Notes: Each column reports the post-period slope of log emissions with respect to the installation-specific loss in free allocation (dose), using the preferred implementation window with installation and year fixed effects and firm-clustered standard errors. The post sample expands from 2021 only to 2021–2022 and then to 2021–2023. The main effects of dose and post are absorbed by fixed effects and omitted due to collinearity.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Controls specification

Across the three specifications in Table 14 the post reform dose response is negative stable and precise. Interpreted per plus ten points in intensity the magnitudes correspond to about minus three point eight percent minus four point four percent and minus three point six percent lower emissions after 2021 respectively.

Table 14: Average post effect (continuous DiD), implementation-only with explicit controls (incl. electricity price)

	(4)	(5)	(6)
Post × Intensity	−0.374*** (0.086)	−0.448*** (0.114)	−0.371*** (0.086)
EUA price			−0.00741*** (0.00083)
ln(GDP volume)	−0.119 (0.393)	−0.226 (0.408)	−1.059** (0.405)
ln(Investment)	−0.084 (0.174)	−0.047 (0.172)	−0.373** (0.158)
ln(Exports goods)	−0.053 (0.326)	0.043 (0.325)	−0.242 (0.307)
ln(Imports goods)	0.021 (0.368)	−0.118 (0.387)	1.368*** (0.264)
ln(Electricity price ≥ 150,000 MWh)	−0.025 (0.036)	−0.039 (0.036)	0.043* (0.025)
Installation FE	Yes	Yes	Yes
Year FE	Yes		
Sector×Year FE		Yes	
Observations	10,482	10,482	10,482
Clusters (firm)	1,512	1,512	1,512
Within R^2	0.0105	0.0107	0.1043

Notes: Dependent variable is ln(emissions). All columns use the implementation-only window (dropping 2019–2020). **(4)** Installation FE + Year FE, controls are ln(GDP volume), ln(Investment), ln(Exports of goods), ln(Imports of goods), and ln(Electricity price, non-household, ≥ 150,000 MWh, excl. taxes) at the country–year level. **(5)** Installation FE + Sector (NACE L1)×Year FE, same controls as (4). **(6)** Installation FE only, includes the same macro controls and EUA price; no time FE so that the year-only EUA series is not collinear with fixed effects. Standard errors in parentheses, clustered at the parent-firm level. “Post × Intensity” is $\text{post}_t \times \text{intensity}_i$. Intensity is scaled in $[0, 1]$, effects per +10 pp equal $0.10 \times \beta$ (percent change $\approx 100(e^{0.10\beta} - 1)\%$). Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample size is smaller here because electricity price coverage limits the country–years to which installations can be matched.

Non parametric bins and continuous at bin means

Table 15: Effect of Free Allocation Loss on $\ln(\text{Emissions})$ by Treatment-Intensity Bins

	(1) A1: Full 2016–2023	(2) A2: Excl. 2019–2020
Bin 2 \times Post	–0.0643* (0.0323)	–0.0804** (0.0374)
Bin 3 \times Post	–0.0684** (0.0328)	–0.0926** (0.0380)
Bin 4 \times Post	–0.1014*** (0.0346)	–0.1395*** (0.0406)
Bin 5 \times Post	–0.2048*** (0.0450)	–0.2540*** (0.0521)
Observations	18,024	13,518
Clusters (firm)	1,614	1,614
R ² (within)	0.0076	0.0111
Root MSE	0.3959	0.4321
Absorbed FE	inst_num, year	inst_num, year
Years in regression	2016–2023	2016–2018, 2021–2023

Notes: Heteroskedasticity-robust SE clustered at firm level in parentheses.
 Bin 1 (lowest treatment intensity) is omitted as the reference group.
 Main effects of bins and of `post_treatment` are collinear with FE and omitted.
 * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 16: Continuous DiD evaluated at bin means of treatment intensity

Bin (quintiles)	Effect at bin mean T		Relative to Bin 1	
	Mean T	Post effect (SE)	Post effect (Bin k – Bin 1)	SE
Bin 1 (lowest T)	0.0515	−0.0163 (0.0033)	0.0000 (reference)	0.0000
Bin 2	0.1442	−0.0456 (0.0093)	−0.0293***	0.0059
Bin 3	0.2419	−0.0765 (0.0155)	−0.0602***	0.0122
Bin 4	0.3742	−0.1183 (0.0240)	−0.1020***	0.0207
Bin 5 (highest T)	0.6740	−0.2131 (0.0432)	−0.1968***	0.0399
Estimator	HDFE: $\ln E_{it} = \alpha_i + \lambda_t + \beta(\text{intensity}_i \times \text{Post}_t) + \varepsilon_{it}$			
$\hat{\beta}$ (SE)	−0.3161 (0.0641)			
Notes	Effects in log points. “Effect at bin mean” equals $\hat{\beta} \times \bar{T}_k$. “Relative to Bin 1” equals $\hat{\beta} \times (\bar{T}_k - \bar{T}_1)$, aligning the baseline with the non-parametric bins table where Bin 1 \times Post is omitted. Standard errors are delta-method: $\text{SE}(\hat{\beta} \cdot a) = a \cdot \text{SE}(\hat{\beta})$, treating \bar{T}_k as fixed. Installation and year fixed effects absorbed; SEs clustered at the firm level. *** $p < 0.01$.			

HAD diagnostics and pre trend checks

Table 17: HAD robustness check (De Chaisemartin–D’Haultfœuille)

	Estimate	SE	95% CI	N	BW	N in BW
Effect_1 (post)	-0.2967	0.2201	[-0.8651, -0.0024]	2,253	0.219	1,048
QUG test: T	1.625 (p = 0.381)					

Notes: Pre years = 2016–2018, post years = 2021–2023. Treatment is set to zero in all pre years and to intensity_{*i*} from 2021 onward. HAD compares treated units to quasi-untreated controls selected by similar pre-trends within a data-driven bandwidth (BW). QUG test assesses comparability of the quasi-untreated group. SEs as reported by the estimator.

Table 18: HAD with pre-trend-matched QUG (post only), 2016-2018 vs. 2021-2023

	Estimate	SE	95% CI	N	BW	N in BW
Effect_1 (post)	-0.2973	0.2216	[-0.8712, -0.0023]	2,252	0.2199	1,051
QUG test: $T = 1.625, p = 0.381$						

Notes: HAD estimator with treatment coded as $D_{it} = \text{intensity}_i \cdot \mathbb{1}\{t \geq 2021\}$, where intensity_{*i*} is the installation-level dose (loss of free allocation). Pre-trend matching keeps post-period QUG units whose 2016–2018 slope is within a data-driven bandwidth (BW) of the median treated pre-trend. Standard errors as reported by the HAD routine.

Table 19: Pre-trend slope regressed on treatment dose intensity_{*i*}

Pre-trend slope (2016–2018)	
Treatment dose intensity _{<i>i</i>}	−0.0154 (0.0163)
Constant	0.0174*** (0.0044)
Observations	2,253
R^2	0.0004

Notes: Each installation’s pre-trend slope is estimated from a regression of $\ln E$ on a linear time index using 2016–2018 data. Table reports OLS of slope_{*i*} on intensity_{*i*} with robust standard errors in parentheses (Stata: `reg slope c.T_i, vce(robust)`). *** $p < 0.01$.

Event study extension

We extend the event study window back to 2015 to align with the Phase III list update. Including 2015 introduces a lead immediately following that update and generates a visible pre period deviation. Consistent with a first year accommodation effect the 2015 coefficient departs from zero and drives the joint pre trend test. We therefore treat 2015 as a transitional year. Our preferred specifications exclude 2015 from estimation which restores stable pre trends and sharpens the post 2021 dose response.

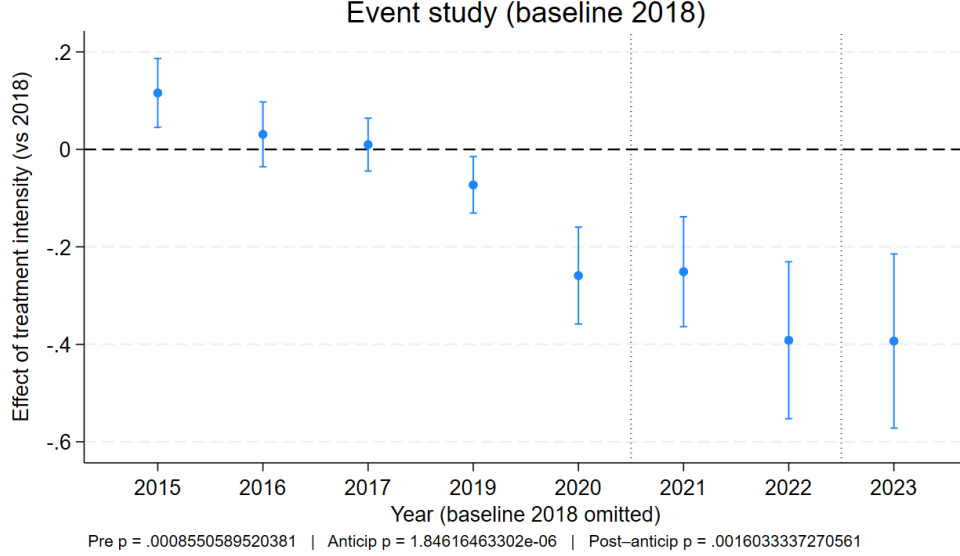


Figure 5: Event study (baseline 2018, 2015-2023). Coefficients are interactions of year dummies with treatment intensity; installation FE, SEs clustered by parent firm.

9.5 Banking construction measures and estimates

Construction of the firm year panel 2015 to 2023

Starting from installation year data we keep firms with at least one Phase III carbon leakage installation in 2015 to 2019. We aggregate verified emissions and free allocation to the firm year level $E_{ft} = \sum_i E_{it}$ and $FA_{ft} = \sum_i FA_{it}$. We compute the annual surplus flow $\Delta_{ft} = FA_{ft} - E_{ft}$. We form the pre 2020 bank as the cumulative surplus up to 2019 capped at zero $\text{bank}_{f,2019}^{\text{capped}} = \max\{0, \sum_{t \leq 2019} \Delta_{ft}\}$. We define the baseline level $E_{\text{base},f}$ as the 2015 to 2019 average of E_{ft} . We construct the firm exposure to the reform TreatExp_f as the portfolio weighted installation dose using 2015 to 2019 emission shares as weights. Banking tests are restricted to firms with positive exposure and a positive pre 2020 bank.

Moderator buffer ratio

We scale the bank by one baseline year of expected shortfall

$$\text{BufferNeed}_f = \frac{\text{bank}_{f,2019}^{\text{capped}}}{\text{TreatExp}_f \cdot E_{\text{base},f}}, \quad \text{for positive } \text{TreatExp}_f.$$

A value of one means the firm can buffer one baseline year of the expected loss from the free allocation cut.

Estimating equations

precise the difference in both table, how each E A and BufferNeed is constructed.

With firm and year fixed effects and $Post_t = \mathbb{1}\{t \geq 2021\}$ we estimate

$$\ln E_{ft} = \alpha_f + \lambda_t + \beta_1 (Post_t \times TreatExp_f) + \beta_2 (Post_t \times BufferNeed_f) + \varepsilon_{ft}, \quad (4)$$

$$\ln E_{ft} = \alpha_f + \lambda_t + \sum_{g \in \{\text{Low, High}\}} \gamma_g (Post_t \times TreatExp_f \times \mathbb{1}\{BufferNeed_f \in g\}) + \varepsilon_{ft}. \quad (5)$$

Table 20: Pre-2020 banking and attenuation of the CL reform effect (firm-year, 2015–2023)

	Dependent variable: $\ln E_{ft}$	
	(1) Continuous ratio-dose	(2) Median split (post slope)
$Post \times TreatExp_f$	−0.293*** (0.070)	
$Post \times BufferNeed_f$	4.3×10^{-6} (8.5×10^{-6})	
Low buffer/need: $Post \times TreatExp_f$		−0.282*** (0.099)
High buffer/need: $Post \times TreatExp_f$		−0.335*** (0.103)
High – Low (difference)		−0.053 (0.143)
Firm FE, Year FE	Yes, Yes	Yes, Yes
Clusters (firms)	1,614	1,614
Observations	12,912	12,912
Within R^2	0.0088	0.0094

Notes: $Post_t = \mathbb{1}\{t \geq 2021\}$. $TreatExp_f$ is the firm’s portfolio-weighted exposure computed from Phase III carbon-leakage incumbents; $BufferNeed_f$ is the pre-2020 bank scaled by one baseline year of the expected shortfall. Column (1) estimates (4); column (2) reports post slopes by median split of $BufferNeed_f$ and their difference. The sample includes 1,614 firms, each operating on average 2.38 installations (range: 1–19) between 2015 and 2023. Robust SEs clustered by firm in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 21: Pre-2020 banking and attenuation of the CL reform effect (firm-year totals, all installations, excl. 2015)

	Dependent variable: $\ln E_{ft}$	
	(1) Continuous ratio-dose	(2) Median split (post slope)
$Post \times TreatExp_f$	-0.264*** (0.060)	
$Post \times BufferNeed_f$	3.16×10^{-6} (8.34×10^{-6})	
Low buffer/need: $Post \times TreatExp_f$		-0.198*** (0.057)
High buffer/need: $Post \times TreatExp_f$		-0.340*** (0.118)
High – Low (difference)		-0.142 (0.124)
Firm FE, Year FE	Yes, Yes	Yes, Yes
Clusters (firms)	1,614	1,614
Observations	12,912	12,912
Within R^2	0.0078	0.0083

Notes: $Post_t = \mathbb{1}\{t \geq 2021\}$. $TreatExp_f$ is the firm's portfolio-weighted exposure computed from Phase III carbon-leakage incumbents. Column (1) estimates (4) with and without the banking moderator $BufferNeed_f$. Column (2) reports post slopes by median split of $BufferNeed_f$ and their difference. The sample includes 1,614 firms, each operating on average 6.68 installations (range: 1–58) between 2016 and 2023. Robust SEs clustered by firm in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

9.6 Accounting variable

Table 22: Firm-level ATE (Post \times dose), preferred spec (bvdid FE; NACE L1 \times year FE)

Outcome	Coef.	Std. err.	t	p	Obs.	Clusters
ln_emissions_firm	-0.273***	0.101	-2.710	0.007	5,954	1,184
ln_empl	-0.088	0.071	-1.240	0.214	5,258	1,140
ln_opre	0.144	0.209	0.690	0.490	5,554	1,113
ln_toas	0.123	0.216	0.570	0.567	5,817	1,141
ln_tfas	-0.101	0.206	-0.490	0.626	5,722	1,129
asinh_ebta	1.685	3.288	0.510	0.609	5,397	1,065
asinh_staf	0.037	0.217	0.170	0.866	5,016	988
asinh_pl	8.764**	4.185	2.090	0.037	5,460	1,074
asinh_plbt	6.841*	4.007	1.710	0.088	5,663	1,114
asinh_margin	-0.255	0.232	-1.100	0.271	5,225	1,038
asinh_turnover	0.061	0.191	0.320	0.748	5,508	1,094
asinh_lev_toas	0.322	0.232	1.390	0.164	5,460	1,074
asinh_cur_lev	0.289	0.237	1.220	0.223	5,663	1,114
asinh_labor_share	-0.085	0.183	-0.460	0.644	4,858	962

Notes: Coefficients are from `reghdfe` with bvdid fixed effects and NACE Level 1 \times year fixed effects; standard errors clustered by bvdid. Variables that are strictly positive are logged; variables that can be zero or negative use the inverse hyperbolic sine. All values rounded to three decimals. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 23: Firm outcomes: construction and economic interpretation

Variable	Construction (this paper)	Economic meaning
ln(Emissions)	$\ln(\sum_{i \in f} E_{it})$	Scale of covered emissions at the BvD entity.
ln(EMPL)	ln(employees)	Real activity and capacity.
ln(OPRE)	ln(operating revenue)	Sales activity near production scale.
ln(TOAS)	ln(total assets)	Capital stock and size.
ln(TFAS)	ln(fixed assets)	Installed capital intensity.
asinh(EBTA)	asinh(EBITDA)	Operating cash-flow proxy with zeros and negatives kept.
asinh(STAF)	asinh(staff costs)	Labor cost bill.
asinh(PL)	asinh(total liabilities)	Leverage and external financing pressure.
asinh(PLBT)	asinh(current liabilities)	Short-term financing or liquidity stress.
asinh(margin)	asinh(EBTA/OPRE)	Operating margin ratio.
asinh(turnover)	asinh(OPRE/TOAS)	Asset turnover, efficiency of the asset base.
asinh(lev_toas)	asinh(PL/TOAS)	Leverage scaled by assets.
asinh(cur_lev)	asinh(PLBT/TOAS)	Short-term leverage scaled by assets.
asinh(labor_share)	asinh(STAF/OPRE)	Labor cost share in revenue.

Notes: $\text{asinh}(x) = \ln(x + \sqrt{x^2 + 1})$ behaves like $\ln(2x)$ for large x and is defined at zero and for negatives, so it keeps balance-sheet and margin observations with losses or zeros. Log transforms are applied only to strictly positive variables.