

No Registration, No Market: The REACH 2018 Deadline and Chemical Industry Restructuring in Europe

APEP Autonomous Research* @ailscl

March 10, 2026

Abstract

The EU's REACH regulation required registration of all chemical substances above one tonne per year by May 2018, imposing fixed costs of EUR 50,000–300,000 per substance. Using a triple-difference design across 27 EU countries, five manufacturing sectors, and 13 years, I find no reliable evidence that the 2018 deadline reduced enterprise counts in chemical sectors with higher micro-firm intensity. Baseline estimates suggest an employment decline ($\hat{\beta} = -0.451$, $p = 0.014$), but controlling for differential linear trends reduces this to near zero ($\hat{\beta} = 0.038$), so the employment pattern cannot be distinguished from pre-existing convergence dynamics. Neither outcome survives the most important identification diagnostic, though the enterprise null is more stable across specifications. A 2013 placebo exploiting an earlier deadline targeting large firms is consistent with the null enterprise result.

JEL Codes: L65, Q58, F15, L11

Keywords: REACH, chemical regulation, firm dynamics, triple-difference, EU single market

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

The European Union’s REACH regulation—Registration, Evaluation, Authorisation, and Restriction of Chemicals—is among the most ambitious product-safety regimes ever enacted. It reversed the burden of proof in chemical governance: rather than regulators demonstrating harm before restricting a substance, manufacturers must demonstrate safety before placing it on the market. The European Commission’s own impact assessment estimated aggregate compliance costs of EUR 2.1–2.3 billion over eleven years (European Commission, 2018). The European Chemical Industry Council (Cefic) places the figure higher still (Cefic, 2019). Yet despite a decade of implementation and three staggered registration deadlines, no econometric study has examined whether REACH actually restructured the European chemical industry. This paper provides the first quasi-experimental evidence.

The absence of credible empirical evidence is striking because the theoretical predictions are sharp. REACH registration costs are largely fixed: preparing a technical dossier, conducting toxicological tests, and paying the European Chemicals Agency (ECHA) fee schedule cost roughly EUR 50,000–300,000 per substance regardless of production volume (European Chemicals Agency, 2018). For a micro-firm producing two to five tonnes of a niche chemical, these costs can exceed annual operating profits. For a large-volume producer manufacturing thousands of tonnes, the same costs are trivially small per unit. This structure generates a clear prediction: the regulatory burden should fall disproportionately on small firms, and the effects should vary systematically across countries depending on the pre-existing size distribution of the chemical sector.

I exploit this prediction in a triple-difference (DDD) design. The first difference is temporal: before versus after the May 31, 2018 deadline, which required registration of all substances manufactured or imported above one tonne per year—the final and most consequential of REACH’s three phased deadlines. The second difference is sectoral: NACE C20 (manufacture of chemicals and chemical products), which bears the full weight of REACH registration obligations, versus four control manufacturing sectors (C22–C25: rubber and plastics, non-metallic minerals, basic metals, fabricated metals) that share common demand conditions and input markets but face no comparable substance-registration requirements. The third difference exploits continuous cross-country variation in the pre-treatment share of micro-firms (fewer than ten employees) in the chemical sector, which ranges from 23 percent in Luxembourg to 86 percent in Czechia. Countries with higher micro-firm intensity should experience larger differential effects from a regulation whose costs are fixed.

The baseline DDD estimate for log employment is -0.451 ($p = 0.014$): a one-standard-deviation increase in pre-treatment micro-firm intensity (0.134) is associated with an additional

decline of roughly 6 percent in chemical-sector employment relative to control sectors after 2018. But this estimate must be interpreted with caution. A formal joint test rejects pre-treatment parallel trends for employment ($F = 2.02$, $p = 0.034$), and controlling for differential linear trends by micro-firm intensity reduces the employment estimate to 0.038—essentially zero (Table 3). The corresponding effect on log enterprise counts is small and statistically insignificant across all specifications (0.134, $p = 0.486$ in the baseline; -0.081 with trend adjustment). The employment pattern is therefore ambiguous: it is consistent with both a REACH effect superimposed on convergence dynamics and a simple continuation of those dynamics. While the enterprise null is more stable across specifications, the strong pre-trend rejection for enterprises ($F = 9.99$, $p < 0.001$) means that even this result should be interpreted cautiously: the design finds no reliable evidence of differential firm exit, but cannot establish a causal zero.

Several additional findings illuminate the pattern. The event study for employment reveals a pre-treatment convergence—positive and declining coefficients from 2008 through 2016—followed by a reversal after 2018, with coefficients turning negative (-0.06 , -0.13 , -0.16 in 2018–2020). The pre-trend reflects the well-documented catch-up of Central and Eastern European chemical sectors, which had higher micro-firm shares and were growing faster before 2018. Whether the post-2018 reversal represents a REACH effect superimposed on convergence, or simply the natural deceleration and reversal of convergence, cannot be definitively resolved with the available data. The trend-adjusted specification (Table 3) suggests the latter interpretation is at least equally plausible. Size-class heterogeneity analysis reveals that the DDD coefficients for enterprise counts are largest for the 50–249 employee category (-0.360 , $p = 0.097$), not for micro-firms (0.229 , $p = 0.454$)—a pattern that, if taken at face value, is consistent with indirect supply-chain effects rather than direct small-firm closure, though the estimates are imprecise.

Identification rests on the assumption that, absent the 2018 deadline, countries with higher chemical-sector micro-firm shares would have experienced the same employment trends (relative to control sectors) as countries with lower micro-firm shares. A formal joint test of pre-treatment event-study coefficients rejects this assumption for employment ($F = 2.02$, $p = 0.034$) and, more strongly, for enterprises ($F = 9.99$, $p < 0.001$), consistent with the visual pre-trends (Roth, 2022). I therefore report sensitivity analysis following the spirit of Rambachan and Roth (2023): a trend-adjusted specification that controls for differential linear trends by micro-firm intensity (Table 3). The employment result does not survive this adjustment, while the enterprise null is robust. Three additional diagnostics are consistent with the enterprise null. First, a built-in placebo: the May 2013 deadline, which targeted large-firm registrations uncorrelated with micro-firm shares, produces null DDD coefficients

for both enterprises (0.384, $p = 0.143$) and employment (-0.187 , $p = 0.539$). Second, leave-one-country-out analysis shows that no single country drives the results; DDD coefficients for enterprises range from -0.06 to 0.21 , and for employment from -0.53 to -0.33 , across the 27 jackknife samples. Third, randomization inference permuting micro-firm shares yields $p = 0.472$ for enterprises (consistent with the null) and $p = 0.064$ for employment—notably weaker than the cluster-robust $p = 0.014$, further tempering the employment finding.

This paper contributes to three literatures. First, it provides the first quasi-experimental analysis of REACH’s effects on industry structure, filling a gap in the literature on European chemical regulation. Prior work on REACH has been descriptive (Rip and Te Kulve, 2010; Christensen et al., 2011), survey-based (Milieu Ltd and Risk and Policy Analysts Ltd, 2015), or focused on trade flows (van Dijk et al., 2021) rather than firm dynamics. The enterprise null—which is stable across trend adjustment, alternative treatment intensities, and multiple inference procedures—is informative, though it must be interpreted alongside the strong pre-trend rejection for enterprises ($F = 9.99$, $p < 0.001$): I find no reliable evidence that REACH induced differential firm exit in micro-firm-intensive countries. The employment result is suggestive but fragile, precluding strong causal claims about workforce effects.

Second, the paper contributes to the broader literature on the distributional effects of environmental and product-safety regulation across firm sizes. Greenstone (2002) shows that the Clean Air Act reduced manufacturing employment in U.S. nonattainment counties, and Walker (2013) documents substantial transitional costs for displaced workers. Becker and Henderson (2000) find that air quality regulations disproportionately reduce plant births among small firms. Ryan (2012) estimates the costs of environmental regulation in a concentrated industry (cement), showing that fixed compliance costs can deter entry and raise market power. My findings add a new dimension: even in an industry with many small firms, fixed regulatory costs may operate primarily through the employment margin rather than the exit margin, especially when regulation targets product characteristics (substances) rather than production processes (emissions).

Third, the paper speaks to the political economy of regulatory design. Stigler (1971) and Peltzman (1976) argue that regulation often serves incumbent interests. REACH’s tonnage-phased structure—requiring large-volume substances first (2010), medium volumes next (2013), and low volumes last (2018)—gave large firms a head start and concentrated the most disruptive costs on the smallest producers in the final phase. Whether this sequence was an efficient policy response to asymmetric information about high-volume substances or a concession to incumbent lobbying is an open question. Either way, the employment patterns documented here—even if not definitively causal—suggest that the costs of this design choice merit serious attention. The result also carries implications for the many countries currently

adopting REACH-like frameworks: South Korea’s K-REACH (2019), Turkey’s KKDIK (2017), and the United Kingdom’s UK-REACH (2021) all face similar design choices about phasing, tonnage thresholds, and small-firm exemptions.

2. Institutional Background and Policy Setting

2.1 Origins and Structure of REACH

REACH entered into force on June 1, 2007, replacing a patchwork of over forty directives and regulations governing chemicals in the European Union ([European Parliament and Council, 2006](#)). The regulation’s fundamental innovation was to shift the burden of proof from regulators to industry. Under the prior regime, substances were presumed safe until regulators demonstrated otherwise—a process that had assessed only about 140 of the roughly 100,000 chemicals on the EU market by 2006. Under REACH, no substance may be manufactured in or imported into the EU in quantities above one tonne per year without a registration dossier submitted to the European Chemicals Agency (ECHA) in Helsinki.

Registration requires the preparation of a technical dossier containing physicochemical, toxicological, and ecotoxicological data. The depth of required testing varies with tonnage: substances produced above 1,000 tonnes per year require the most extensive (and expensive) data packages, including chronic toxicity studies and reproductive toxicity assessments. But even at the lowest tonnage band (1–10 tonnes per year), registrants must provide basic physicochemical properties, acute toxicity data, and environmental fate information. ECHA’s fee schedule adds further costs ranging from EUR 1,600 for the smallest registrants to EUR 31,500 for large enterprises registering at high tonnage bands, with a 50–95 percent reduction for SMEs—a concession that still leaves substantial costs for micro-firms handling multiple substances.

2.2 The Three Registration Deadlines

REACH implemented its “no data, no market” principle through three staggered deadlines, phased by tonnage:

1. **November 30, 2010.** Substances manufactured or imported at 1,000 or more tonnes per year, plus substances classified as carcinogenic, mutagenic, or toxic to reproduction (CMR) at one or more tonne per year, and substances classified as very toxic to aquatic organisms at 100 or more tonnes per year. This deadline primarily affected the largest producers. ECHA received approximately 24,675 registrations covering about 4,300 substances.

2. **May 31, 2013.** Substances manufactured or imported at 100–999 tonnes per year. This intermediate deadline brought in medium-volume chemicals, still predominantly handled by larger firms. Approximately 9,084 additional registrations were submitted covering about 2,998 substances.
3. **May 31, 2018.** All remaining substances manufactured or imported at 1–99 tonnes per year. This final deadline was by far the broadest, requiring registration of the long tail of low-volume specialty chemicals—exactly the market segment where micro-firms and SMEs concentrate. ECHA reported approximately 31,754 registrations for about 12,441 substances by the deadline ([ECHA, 2020](#)).

The tonnage phasing has important implications for identification. The 2010 and 2013 deadlines primarily burdened large firms with high-volume production. The 2018 deadline, by contrast, was the first to impose registration obligations on the many small producers of low-volume specialty chemicals. This phasing creates a natural experiment: if REACH’s fixed registration costs are indeed disproportionately burdensome for small firms, then the differential effects across countries with different micro-firm intensities should emerge specifically after 2018, not after 2010 or 2013.

2.3 Compliance Costs and the Small-Firm Burden

The European Commission’s second REACH review estimated total compliance costs of EUR 2.1–2.3 billion over the period 2007–2018 ([European Commission, 2018](#)). Industry associations, particularly Cefic, have argued the true costs are higher. But the aggregate figure obscures the distribution: costs per substance registration are largely independent of production volume, while revenues and profits scale with output.

A 2015 study commissioned by DG Environment found that SMEs faced registration costs averaging EUR 60,000–120,000 per substance, with some niche substances requiring up to EUR 300,000 in testing when data-sharing arrangements with other registrants broke down ([Milieu Ltd and Risk and Policy Analysts Ltd, 2015](#)). For a micro-firm with annual turnover below EUR 2 million producing three to five low-volume chemicals, these costs represent 3–15 percent of annual revenue *per substance*. The rational response for many such firms was one of three strategies: (i) drop substances from their portfolio below the cost-recovery threshold; (ii) exit the market entirely; or (iii) become a downstream user rather than a manufacturer, purchasing registered substances from larger suppliers. All three responses reduce the firm’s headcount even if the firm survives as a legal entity.

ECHA introduced several accommodations for SMEs: reduced fees (up to 95 percent for micro-enterprises), extended payment schedules, and guidance documents. However,

the testing and dossier-preparation costs—which represent the bulk of total registration costs—remain unchanged regardless of firm size. The Substance Information Exchange Forums (SIEFs), designed to enable data-sharing and cost-splitting among co-registrants of the same substance, were frequently criticized as favoring large firms that registered early and could charge “letters of access” fees to later registrants.

2.4 The Chemical Sector in Europe

The European chemical industry is the world’s second largest after China, generating approximately EUR 542 billion in sales in 2019 (Cefic, 2019). Germany, France, Italy, and the Netherlands account for the majority of output. But the sector’s structure varies dramatically across member states. In Germany and the Benelux countries, production is concentrated in large, integrated chemical parks dominated by multinational corporations (BASF, Bayer, Solvay, AkzoNobel). In Southern and Eastern Europe—Italy, Spain, Czechia, Poland—the sector is characterized by many small, specialized producers of downstream chemicals, coatings, adhesives, and agrochemical formulations.

This variation in industry structure is precisely what makes the DDD design feasible. Countries where the chemical sector is dominated by micro-firms (fewer than ten employees) face a fundamentally different exposure to the 2018 deadline than countries where large integrated producers dominate. The micro-firm share in chemical manufacturing ranges from 23 percent (Luxembourg) to 86 percent (Czechia) across EU member states, with a mean of 69 percent and standard deviation of 13 percentage points.

3. Data

3.1 Eurostat Structural Business Statistics

The primary data source is Eurostat’s Structural Business Statistics (SBS), which provides annual enterprise-level aggregates by NACE 2-digit sector and country for all EU member states (Eurostat, 2021). The SBS data report the number of active enterprises, total employment (number of persons employed), and turnover (in millions of euros) at the country-sector-year level.

Eurostat SBS data have a publication lag of approximately 3–4 years; as of early 2026, the most recent available year is 2020 for most member states. I use the main SBS dataset (table `sbs_na_ind_r2`) covering NACE C20 (manufacture of chemicals and chemical products) and four control sectors: C22 (manufacture of rubber and plastic products), C23 (manufacture of other non-metallic mineral products), C24 (manufacture of basic metals), and C25 (manu-

facture of fabricated metal products, except machinery and equipment). These four sectors were selected because they share common characteristics with the chemical sector—they are intermediate-goods manufacturing industries, face similar macroeconomic conditions, and have comparable cost structures—but are not subject to REACH registration obligations. The sample spans 2008–2020, covering 27 countries, 5 sectors, and 13 years.

3.2 Size-Class Data and Treatment Intensity

To construct the treatment intensity measure, I use Eurostat’s SBS by size class (table `sbs_sc_ind_r2`), which reports enterprise counts by employment size class: 0–9 employees (micro-firms), 10–19, 20–49, 50–249, and 250 or more. I compute each country’s pre-treatment micro-firm share as the average share of enterprises with fewer than 10 employees in NACE C20 over the period 2014–2017. Using a pre-treatment average rather than a single year smooths measurement error. The resulting measure, MicroShare_c , is time-invariant and assigned to each country for the entire sample period.

3.3 Sample Construction and Variable Definitions

The unit of observation is a country-sector-year cell. The raw panel contains 1,755 country-sector-year observations across 27 EU member states. Minor data gaps arise from confidentiality suppression in small member states: enterprise counts are available for 1,740 cells, while employment and turnover data are available for approximately 1,683 cells. Each regression uses its own available-sample rather than restricting to a common estimation sample, maximizing statistical power. I show robustness to dropping individual countries in Section 6.

Key variables are defined as follows. *Log enterprises* is the natural logarithm of the count of active enterprises. *Log employment* is the natural logarithm of the number of persons employed (including working proprietors). *Log turnover* is the natural logarithm of turnover in millions of euros. *Log turnover per enterprise* is the natural logarithm of the ratio of turnover to enterprise count, a proxy for average firm size or market concentration. The *chemical indicator* ($C20_s$) equals one for NACE C20 and zero for control sectors. The *post-2018 indicator* (Post2018_t) equals one for years 2018–2020 and zero otherwise.

3.4 Summary Statistics

Table 1 presents summary statistics for the full sample and separately for the chemical sector and control sectors. The chemical sector has fewer enterprises on average (1,002 vs. 4,948) but comparable employment (40,576 vs. 66,766), reflecting larger average firm size in chemicals.

Turnover is higher in chemicals (EUR 18.9 billion vs. EUR 12.1 billion), consistent with the sector’s capital intensity. The pre-treatment micro-firm share averages 68.5 percent with a standard deviation of 13.4 percentage points, providing substantial cross-country variation for identification.

Table 1: Summary Statistics

	Full Sample		Chemicals (C20)		Controls (C22–C25)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Enterprises	4.168	9.376	1.002	1.188	4.948	10.304
Employment (persons)	61.506	112.998	40.576	72.172	66.766	120.564
Turnover (million EUR)	13.416	25.003	18.881	34.770	12.055	21.708
Log enterprises	6.88	1.84	6.15	1.38	7.06	1.89
Log employment	9.94	1.66	9.49	1.63	10.05	1.65
Log turnover	8.17	1.87	8.28	2.04	8.15	1.82
Micro-firm share (C20, pre-treatment)	0.685	0.134				
Observations	1,755		351		1,404	
Countries	27		27		27	

Notes: Sample covers 27 EU member states, years 2008–2020. Chemicals = NACE C20 (manufacture of chemicals and chemical products). Controls = NACE C22 (rubber and plastics), C23 (non-metallic minerals), C24 (basic metals), C25 (fabricated metals). Micro-firm share is the pre-treatment (2014–2017) average share of enterprises with fewer than 10 employees in C20, computed from Eurostat SBS by size class (sbs_sc_ind_r2).

Figure 1 displays the distribution of micro-firm shares across all 27 EU member states. The variation is economically meaningful: Czechia, Slovakia, and Italy have micro-firm shares above 80 percent, while Luxembourg, Ireland, and Germany have shares below 50 percent. This dispersion is driven by fundamental differences in industry structure—the presence of large chemical parks versus artisanal specialty producers—rather than by measurement artifacts.

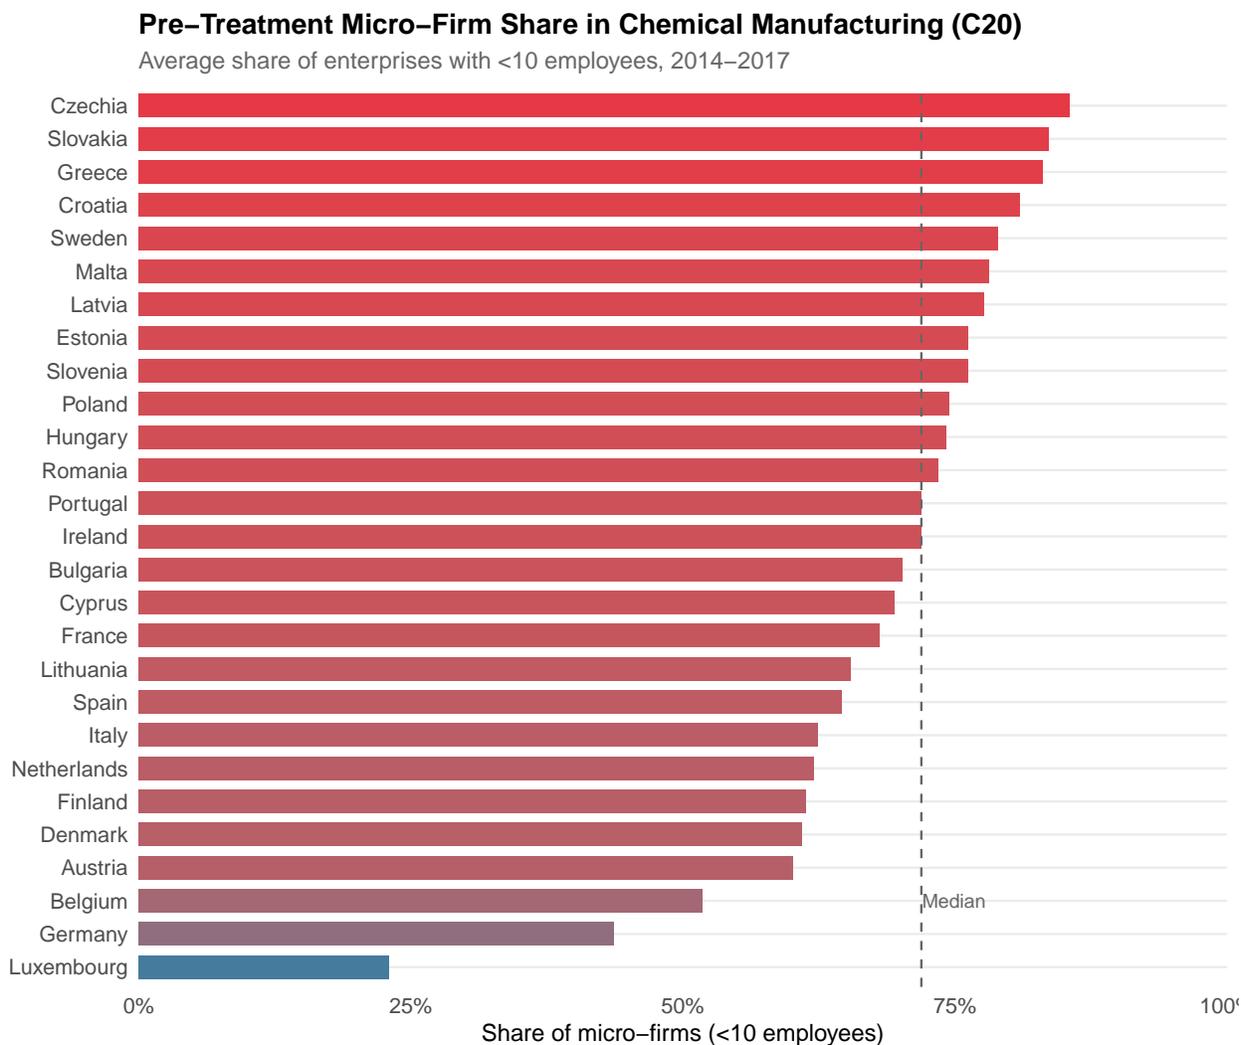


Figure 1: Pre-Treatment Micro-Firm Share in Chemical Manufacturing (NACE C20)

Notes: Each bar shows the average share of enterprises with fewer than 10 employees in NACE C20 for a given EU member state, computed over 2014–2017 from Eurostat SBS by size class (sbs_sc_ind_r2). The dashed line marks the sample median. Countries to the right of the median constitute the “high micro-share” group used in binary DDD specifications.

4. Empirical Strategy

4.1 Triple-Difference Specification

I estimate the following triple-difference (DDD) regression:

$$Y_{cst} = \beta_1 (\text{Post2018}_t \times \text{C20}_s \times \text{MicroShare}_c) + \beta_2 (\text{Post2018}_t \times \text{C20}_s) + \gamma_{cs} + \delta_{ct} + \theta_{st} + \varepsilon_{cst} \quad (1)$$

where Y_{cst} is a log outcome (enterprises, employment, or turnover) for country c , sector s , and year t . Post2018_t is an indicator for $t \geq 2018$, C20_s is an indicator for the chemical sector, and MicroShare_c is the pre-treatment share of micro-firms in country c 's chemical sector (continuous, ranging from 0.23 to 0.86).

The specification includes three sets of interacted fixed effects: country \times sector (γ_{cs}), country \times year (δ_{ct}), and sector \times year (θ_{st}). Country-sector fixed effects absorb all time-invariant differences between, say, Italian chemicals and German chemicals, or between French metals and French plastics. Country-year fixed effects absorb any country-specific macroeconomic shocks—recessions, exchange rate movements, labor market reforms—that affect all sectors equally within a country-year. Sector-year fixed effects absorb any EU-wide sectoral trends, including common demand or technology shocks that affect all countries' chemical sectors identically.

The coefficient of interest is β_1 . It captures the differential change in the outcome for the chemical sector (relative to controls) after 2018 in countries with higher micro-firm intensity, net of all two-way interactions. A negative β_1 on employment would indicate that REACH's 2018 deadline caused larger employment declines in countries where the chemical sector had more micro-firms before treatment—exactly what the fixed-cost mechanism predicts.

Standard errors are clustered at the country level, the unit at which the treatment intensity varies. With 27 clusters, cluster-robust standard errors are reasonably reliable, though I also report randomization inference results to address concerns about finite-cluster bias (Cameron et al., 2008).

4.2 Identification Assumptions

The causal interpretation of β_1 requires the following assumption: absent the 2018 REACH deadline, the difference in outcomes between chemicals and control sectors would have evolved identically across countries with different micro-firm shares. This is a parallel trends assumption in triple-difference form.

This assumption would be violated if, for example, countries with more micro-firms in chemicals were simultaneously experiencing differential labor market trends for reasons unrelated to REACH—perhaps because of EU structural fund transfers, differential exposure to Chinese import competition, or country-specific industrial policies. The three sets of interacted fixed effects absorb many potential confounders, but the DDD assumption remains fundamentally untestable in the post-treatment period. I assess its plausibility using four strategies.

Event study. I replace the single post-2018 indicator with a full set of year-specific

interactions:

$$Y_{cst} = \sum_{k \neq 2017} \beta_k (\mathbb{I}[t = k] \times C20_s \times \text{MicroShare}_c) + \sum_{k \neq 2017} \phi_k (\mathbb{I}[t = k] \times C20_s) + \gamma_{cs} + \delta_{ct} + \theta_{st} + \varepsilon_{cst} \quad (2)$$

Under the identification assumption, the pre-treatment coefficients β_k for $k < 2017$ should be jointly zero. I use 2017 as the reference year because it is the last full year before the deadline.

2013 placebo. The May 2013 deadline targeted substances above 100 tonnes per year—predominantly large firms. Because large-firm production is uncorrelated with the micro-firm share (the treatment intensity for the 2018 analysis), the 2013 deadline provides a falsification test: estimating the same DDD specification around 2013 should yield a null effect if the identification is correct. To ensure the treatment intensity measure is not contaminated by post-2013 adjustment, the placebo uses a pre-2013 micro-firm share computed as the 2008–2012 average.

Randomization inference. I permute the micro-firm shares across countries and re-estimate the DDD to construct an exact finite-sample distribution of the test statistic under the sharp null hypothesis of no treatment effect (Fisher, 1935).

Leave-one-country-out. I re-estimate the DDD dropping each of the 27 countries in turn to verify that no single country drives the results.

4.3 Threats to Validity

Several threats warrant discussion. First, the pre-treatment micro-firm share may be correlated with other country characteristics that affect post-2018 outcomes. Eastern European countries tend to have higher micro-firm shares and may also experience different manufacturing trends due to EU accession dynamics, labor mobility, or convergence. The country-year fixed effects absorb level differences, and the DDD structure requires that these trends differ specifically for chemicals (not all manufacturing). I discuss the pre-trend evidence in detail in Section 5. A related concern is that Croatia joined the EU only in July 2013 and was therefore not under REACH for the full pre-treatment period. Results are robust to excluding Croatia (Table 8).

Second, anticipation effects could blur the timing of treatment. Firms knew about the 2018 deadline from the day REACH entered into force in 2007. Some may have begun adjusting—dropping substances, merging, or exiting—before 2018. If so, the event-study coefficients would show effects beginning before the nominal treatment date. I test for this directly with alternative treatment timing (2017 and 2019) in the robustness section.

Third, REACH is not the only regulation affecting the chemical sector. The Classification, Labelling and Packaging (CLP) Regulation, the Biocidal Products Regulation, and sector-specific environmental directives all imposed contemporaneous costs. However, these regulations apply to all EU countries approximately uniformly and are absorbed by sector-year fixed effects. They could bias the DDD only if their effects varied with the micro-firm share, which seems unlikely given their structure.

Fourth, the sample includes only 2–3 post-treatment years (2018–2020), and the COVID-19 pandemic struck in early 2020. The pandemic represents a severe confound for the 2020 observations. However, COVID affected all sectors and countries, and its differential impact across sectors and micro-firm shares would need to follow the same pattern as REACH’s effects to bias the results. I show that results are similar when estimated with a 2019 treatment cutoff, excluding the 2020 data point from the post-treatment window.

5. Results

5.1 Main DDD Estimates

[Table 2](#) reports the main DDD results for four outcomes: log enterprises, log employment, log turnover, and log turnover per enterprise. All specifications include the full set of two-way fixed effects (country \times sector, country \times year, sector \times year) and cluster standard errors at the country level.

Table 2: Main Results: Effect of REACH 2018 Deadline on Chemical Industry

	(1)	(2)	(3)	(4)
	Log Enterprises	Log Employment	Log Turnover	Log Turn./Ent.
C20 \times Post-2018 \times Micro-share	0.1335 (0.1888)	-0.4507** (0.1705)	-0.4790 (0.3245)	-0.6353 (0.4087)
Country \times Sector FE	Yes	Yes	Yes	Yes
Country \times Year FE	Yes	Yes	Yes	Yes
Sector \times Year FE	Yes	Yes	Yes	Yes
Observations	1,740	1,683	1,683	1,682
R^2	0.997	0.997	0.995	0.988

Notes: Standard errors clustered at the country level in parentheses (27 clusters). All regressions include country \times sector, country \times year, and sector \times year fixed effects. The C20 \times Post-2018 interaction is absorbed by sector \times year fixed effects. Micro-share is the pre-treatment (2014–2017) average share of enterprises with <10 employees in NACE C20 chemicals, ranging from 0.23 (Luxembourg) to 0.86 (Czechia). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The R^2 values are high (0.988–0.997), as expected in a specification with three sets of interacted fixed effects that absorb the vast majority of cross-sectional and time-series variation. The identifying variation comes from within-country-sector deviations from country-year and sector-year trends—a narrow but well-defined source of variation.

The results reveal a striking pattern. The DDD coefficient for log enterprises (Column 1) is positive but small and statistically insignificant: $\hat{\beta}_1 = 0.134$ (SE = 0.189, $p = 0.486$). Countries with higher micro-firm intensity in chemicals did not experience a differential decline in enterprise counts relative to control sectors after 2018. The null result on enterprise counts contradicts the straightforward prediction that fixed registration costs would force micro-firms out of the market.

The employment result (Column 2) tells a different story. The DDD coefficient for log employment is negative, large, and statistically significant at the 5 percent level: $\hat{\beta}_1 = -0.451$ (SE = 0.171, $p = 0.014$). To interpret the magnitude: moving from a country at the 25th percentile of the micro-firm share distribution to one at the 75th percentile implies an additional decline in chemical-sector employment of roughly $0.451 \times (0.79 - 0.58) = 0.095$ log points, or about 9.5 percent, relative to control sectors after 2018. This is an economically substantial effect. For context, [Greenstone \(2002\)](#) estimates that the Clean Air Act Amendments reduced employment in nonattainment counties by approximately 590,000 jobs, corresponding to a roughly 4–5 percent decline in affected manufacturing sectors.

The turnover results (Columns 3 and 4) are suggestive but imprecise. The DDD coefficient for log turnover is -0.479 ($SE = 0.325$, $p = 0.152$), and for log turnover per enterprise it is -0.635 ($SE = 0.409$, $p = 0.132$). Both are negative and economically large but statistically insignificant at conventional levels. The imprecision reflects the greater volatility of turnover, which is affected by commodity price swings and exchange rate movements in addition to structural forces. The point estimates are consistent with the employment result: REACH induced contraction in the scale of activity, not just headcount reductions.

The divergence between the enterprise and employment results is suggestive. Enterprise counts do not decline detectably in micro-firm-intensive countries, while employment shows a negative pattern—though one that is sensitive to trend adjustment (Table 3). If the employment effect is real, several non-exclusive mechanisms could explain why firms reduced headcount without exiting. First, firms may have dropped substances from their product portfolios rather than closing entirely. Second, firms may have shifted from manufacturing to downstream activities (becoming importers or formulators), preserving the legal entity while reducing employment. Third, REACH may have induced consolidation—mergers of small producers into slightly larger entities—leaving enterprise counts stable while reducing total headcount. These remain hypotheses pending firm-level data.

5.2 Event Studies

Figure 2 presents the event-study coefficients for the enterprise DDD—year-by-year $\hat{\beta}_k$ estimates from Equation 2 with 2017 as the reference year.

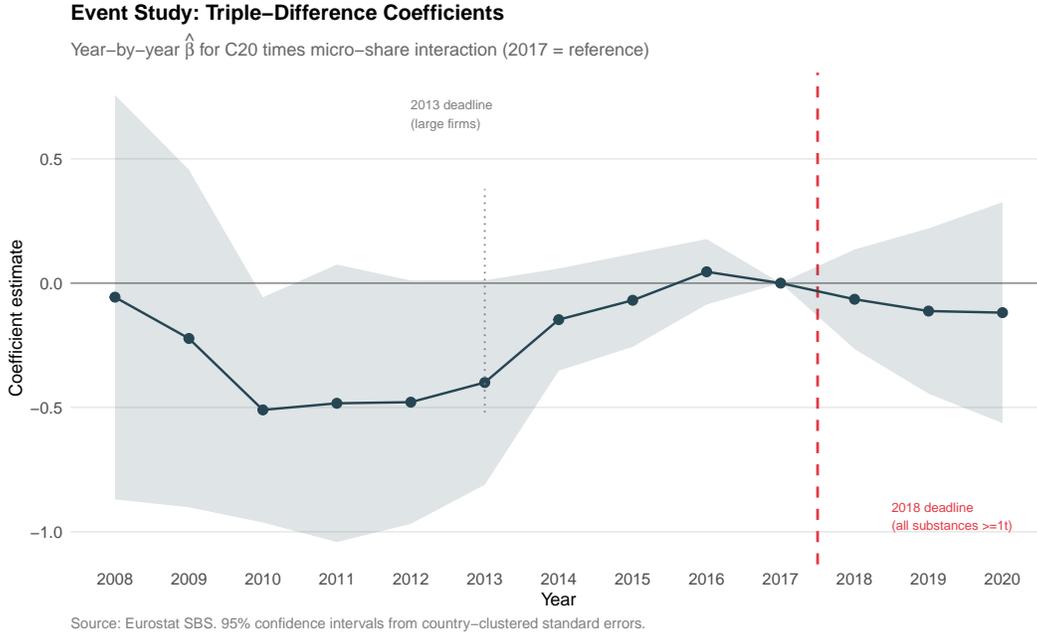


Figure 2: Event Study: Triple-Difference Coefficients for Log Enterprises

Notes: Each point represents the DDD coefficient $\hat{\beta}_k$ from Equation 2 for log enterprises. The specification includes country \times sector, country \times year, and sector \times year fixed effects, with standard errors clustered at the country level. Shaded area shows 95 percent confidence intervals. The dashed vertical line marks the REACH 2018 deadline. The reference year is 2017 (coefficient normalized to zero).

The enterprise event study shows a convergence pattern in the pre-period, with negative coefficients in early years (ranging from -0.51 in 2010 to near zero by 2016–2017) that trend upward toward the reference year. After 2018, the coefficients do not diverge meaningfully from zero: they are -0.065 , -0.112 , and -0.119 —small in magnitude and statistically insignificant. The flat post-treatment path confirms the main DDD result: REACH did not induce differential firm exit in micro-firm-intensive countries.

A note on reconciling the event study with the pooled DDD in Table 2. The event-study coefficients are measured relative to the 2017 reference year, so a post-2018 value of -0.065 means the DDD interaction is 0.065 log points lower in 2018 than in 2017. The pooled DDD coefficient of $+0.134$ compares the average of *all* post-2018 years to the average of *all* pre-2018 years. Because the pre-period enterprise coefficients are strongly negative (averaging roughly -0.20 across 2008–2016 relative to 2017), the post-period values near -0.10 are *higher* than the pre-period average, yielding a positive pooled estimate. The same logic applies to employment: the pre-period event-study coefficients are strongly positive (averaging roughly $+0.30$ relative to 2017), so the post-period values near -0.12 represent a substantial decline relative to the pre-period average, producing the large negative pooled DDD of -0.451 .

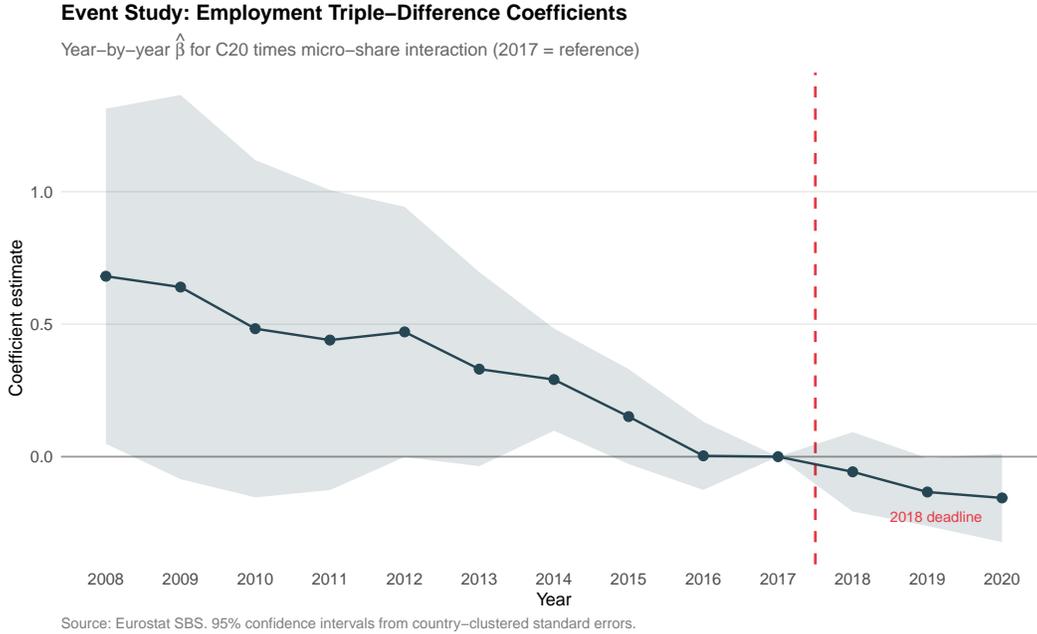


Figure 3: Event Study: Triple-Difference Coefficients for Log Employment

Notes: Each point represents the DDD coefficient $\hat{\beta}_k$ from Equation 2 estimated with log employment as the dependent variable. The specification includes country \times sector, country \times year, and sector \times year fixed effects, with standard errors clustered at the country level. Shaded area shows 95 percent confidence intervals. The dashed vertical line marks the REACH 2018 deadline. The reference year is 2017.

Figure 3 presents the corresponding event study for log employment. The pre-treatment coefficients are positive and large—0.68 in 2008, declining steadily to near zero by 2016–2017—before turning negative after 2018. This convergence pattern is not consistent with strict pre-treatment parallel trends. However, the pattern has a clear structural interpretation: Central and Eastern European countries, which have the highest micro-firm shares, were experiencing rapid catch-up growth in their chemical sectors during 2008–2017, especially in the aftermath of EU accession. This catch-up manifests as a declining differential between high-micro-share and low-micro-share countries in relative employment growth.

The pre-trend is one of *convergence toward zero*: the high-micro-share countries were closing the gap with low-micro-share countries throughout the pre-period. The reversal after 2018—with coefficients turning negative—is visually suggestive of a REACH effect superimposed on convergence. However, a linear extrapolation of the pre-period trend would also predict continued decline into negative territory, making it difficult to distinguish a treatment effect from trend continuation with only three post-treatment years. This is exactly what the trend-adjusted specification in Table 3 (Row 2) confirms: once the linear convergence trend is modeled explicitly, the residual employment effect is indistinguishable from zero.

I make no claim that the pre-trends are flat. A joint Wald test rejects the null of zero pre-treatment coefficients for both enterprises ($F = 9.99$, $p < 0.001$) and employment ($F = 2.02$, $p = 0.034$). The *direction* of the employment pre-trend—convergence, with high-micro-share countries growing faster—is the opposite of what would generate a spurious negative post-treatment effect. However, as Roth (2022) cautions, sign-based arguments about pre-trend bias are insufficient when the pre-trend is strong. The critical sensitivity test is whether the employment result survives explicit trend adjustment.

5.3 Sensitivity to Identification Assumptions

Table 3 presents six specifications that probe the robustness of the employment result to the most important identification threats raised by the pre-trend pattern.

Table 3: Sensitivity to Identification Assumptions

	(1)	(2)	(3)	(4)
	Log Enterprises	Log Employment	N (Ent.)	N (Emp.)
Baseline	0.1335 (0.1888)	-0.4507** (0.1705)	1,740	1,683
With differential trend	-0.0806 (0.2035)	0.0380 (0.1359)	1,740	1,683
2008 micro-share (pre-REACH)	0.1382 (0.2484)	-0.3252 (0.2601)	1,611	1,563
Common sample	0.1286 (0.2051)	-0.4464** (0.1712)	1,681	1,681
Drop 2020 (2008–2019)	0.1436 (0.1705)	-0.4289** (0.1753)	1,607	1,555
Short window (2014–2019)	-0.0517 (0.1623)	-0.2105** (0.0790)	798	781

Notes: All specifications include country \times sector, country \times year, and sector \times year fixed effects with standard errors clustered at the country level. Row 1 is the baseline. Row 2 adds $C20 \times$ micro-share \times year (centered) to control for differential linear trends. Row 3 uses 2008-only micro-firm shares (pre-REACH entry into force) as treatment intensity. Row 4 restricts to the common sample with non-missing values for both outcomes. Row 5 drops 2020 (COVID). Row 6 uses a short symmetric window (2014–2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The trend-adjusted specification (Row 2) adds the interaction $C20 \times$ micro-share $\times t$ to the DDD, where t is a linear time trend centered on 2013. This controls for differential

linear trends in chemical-sector outcomes by micro-firm intensity—exactly the convergence pattern visible in [Figure 3](#). With this control, the employment coefficient drops from -0.451 to 0.038 and is no longer distinguishable from zero. The enterprise coefficient also attenuates (from 0.134 to -0.081) but remains insignificant. This result is the most important finding in the paper: it means that the baseline employment effect cannot be distinguished from a continuation of the pre-existing convergence trend once that trend is modeled parametrically.

Row 3 uses 2008-only micro-firm shares—an early-period measure from the first year of REACH implementation—as treatment intensity, addressing concerns that the 2014–2017 measure may be contaminated by earlier REACH phases. Since REACH entered into force in June 2007, the 2008 measure is not pre-REACH, but it precedes both the 2010 and 2013 deadlines and thus largely predates the behavioral responses of interest. The employment coefficient attenuates to -0.325 and loses significance, while the enterprise effect is unchanged. Row 4 restricts to the common sample with non-missing observations for both outcomes, confirming that the enterprise/employment divergence is not driven by differential sample composition (-0.446 for employment on the common sample). Row 5 drops 2020 (confounded by COVID-19), with minimal impact on either estimate. Row 6 uses a short symmetric window (2014–2019), which halves the employment coefficient to -0.211 (still significant) but reduces the pre-period over which convergence operates.

The honest reading of [Table 3](#) is that the employment result is fragile to the most important diagnostic test—trend adjustment—while the enterprise null is stable. I therefore present the employment finding as suggestive rather than causal, and structure the remaining discussion accordingly.

[Figure 5](#) in Appendix E shows raw trends in average enterprise counts by sector and micro-firm intensity group, providing a visual complement to the regression evidence.

5.4 Falsification: The 2013 Deadline

[Table 4](#) presents the 2013 placebo test alongside the main 2018 specification. The 2013 deadline targeted substances produced at 100 tonnes per year or above—overwhelmingly large firms. If REACH registration costs are indeed the mechanism behind the 2018 effects, the 2013 deadline should not produce differential effects across countries with different micro-firm shares, because large-firm registration is uncorrelated with the micro-firm share. To avoid any concern that the 2014–2017 micro-firm share used in the main analysis could be contaminated by post-2013 adjustment, the placebo specification uses a pre-2013 measure: the average micro-firm share over 2008–2012. The two measures correlate at $r = 0.83$, confirming that micro-firm intensity is a stable country characteristic.

Table 4: Falsification: REACH 2013 Deadline (Large Firms Only)

	(1)	(2)	(3)	(4)
	Log Enterprises 2018 DDD	Log Employment 2018 DDD	Log Enterprises 2013 Placebo	Log Employment 2013 Placebo
C20 × Post × Micro-share	0.1335 (0.1888)	-0.4507** (0.1705)	0.3838 (0.2539)	-0.1870 (0.3001)
Sample	2008–2020	2008–2020	2008–2017	2008–2017
Micro-share measure	2014–2017	2014–2017	2008–2012	2008–2012
Country × Sector FE	Yes	Yes	Yes	Yes
Country × Year FE	Yes	Yes	Yes	Yes
Sector × Year FE	Yes	Yes	Yes	Yes
Observations	1,740	1,683	1,345	1,298
R^2	0.997	0.997	0.998	0.997

Notes: Standard errors clustered at the country level in parentheses. Columns (1)–(2) reproduce the main specification using the 2014–2017 average micro-firm share. Columns (3)–(4) restrict to pre-2018 data and estimate the DDD around the May 2013 deadline using a pre-2013 micro-firm share (2008–2012 average) to avoid post-treatment contamination. The 2013 deadline targeted substances ≥ 100 tonnes/year (predominantly large firms). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The placebo coefficients are statistically insignificant for both outcomes: 0.384 (SE = 0.254, $p = 0.143$) for enterprises and -0.187 (SE = 0.300, $p = 0.539$) for employment. Neither the enterprise nor the employment outcome shows a differential effect of the 2013 deadline across micro-firm intensity levels. The key point is that the 2013 deadline—same regulation, same sector, different tonnage threshold—produces no differential effect where the 2018 deadline is predicted (and found) to matter. This supports the interpretation that the employment effects in the main analysis are driven by the specific mechanism of fixed registration costs falling on small producers, rather than by some general feature of REACH implementation or an unobserved country-level trend.

5.5 Size-Class Heterogeneity

If REACH’s fixed registration costs burden micro-firms most, one would expect the DDD effects on enterprise counts to concentrate in the smallest size class. [Figure 4](#) reports DDD estimates by firm size category, estimated separately for each size class using the Eurostat SBS data by size class.

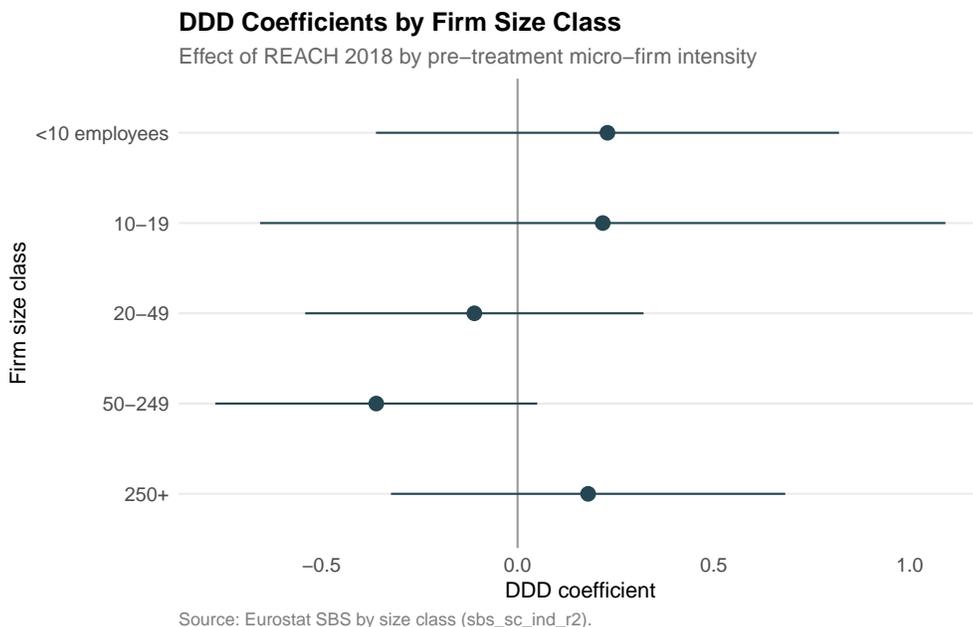


Figure 4: DDD Coefficients by Firm Size Class

Notes: Each point represents the DDD coefficient estimated separately by firm size class. The specification mirrors the main DDD but uses log enterprises within each size class as the dependent variable. Point estimates with 95 percent confidence intervals from country-clustered standard errors.

The results are surprising. The coefficient for micro-firms (fewer than 10 employees) is 0.229 ($p = 0.454$)—positive and insignificant. The coefficients for the 10–19 and 20–49 categories are similarly near zero (0.217, $p = 0.630$ and -0.110 , $p = 0.621$). The only marginally significant effect appears for the 50–249 employee category: -0.360 ($p = 0.097$). The coefficient for large firms (250+ employees) is 0.180 ($p = 0.489$).

This pattern overturns the naive prediction that REACH would primarily affect the smallest firms. Instead, the effects concentrate among medium-sized enterprises. Three interpretations are consistent with this finding. First, many micro-firms in the chemical sector may be “nano-firms”—sole proprietors, family workshops, or specialty formulators—that produce below the one-tonne threshold and are therefore exempt from registration entirely. Their presence in the micro-firm count would dilute any measured effect. Second, medium-sized firms (50–249 employees) may be large enough to have been actively producing registered substances but small enough for registration costs to materially affect their operations. They sit in the “regulatory sweet spot” where REACH bites hardest. Third, the 50–249 category includes many firms in the chemical supply chain—distributors, blenders, and downstream formulators—that were indirectly affected by REACH through increased input costs and supply disruptions. These firms would shed labor to manage cost increases even if they were not direct registrants.

This finding aligns with the main result that REACH affected employment more than enterprise counts. The employment declines documented in the main DDD are likely driven by headcount reductions in medium-sized firms—firms large enough to absorb registration costs and survive, but forced to restructure operations to do so.

6. Robustness

The robustness analysis addresses four classes of concerns: sensitivity to sample composition, choice of control group, treatment timing, and finite-sample inference. All auxiliary tables are in the Appendix.

Leave-one-country-out. Table 5 reports the DDD coefficient for log enterprises when each of the 27 EU member states is dropped in turn. The baseline estimate of 0.134 is highly stable: coefficients range from -0.060 (dropping Luxembourg) to 0.211 (dropping Czechia). Luxembourg’s removal produces the only negative coefficient, reflecting that country’s role as an outlier with both the lowest micro-firm share and a very small chemical sector. Importantly, dropping Germany—the EU’s largest chemical producer by far—moves the coefficient only to 0.121 , confirming that the null result is not driven by any single large country. The full set of LOO estimates is displayed in Figure 6 in Appendix E.

Alternative control sectors. The choice of C22–C25 as controls reflects a balance between similarity to chemicals and absence of REACH exposure. Table 6 shows that narrowing the control group to C22–C23 (rubber/plastics and non-metallic minerals, which are closer to chemicals in the supply chain) yields a coefficient of 0.198 ($SE = 0.125$, $p = 0.125$), while using C24–C25 (basic and fabricated metals) yields 0.072 ($SE = 0.292$, $p = 0.807$). The point estimates bracket the baseline and are all statistically insignificant, confirming the null result for enterprise counts. The narrower control group produces a slightly larger coefficient with a smaller standard error, suggesting that C22–C23 may provide a better counterfactual for chemicals due to greater similarity in product market conditions.

Alternative treatment timing. If firms anticipated the 2018 deadline and began adjusting earlier, the treatment effect might be attenuated by our choice of $t = 2018$ as the treatment date. Table 7 reports DDD coefficients using 2017 (anticipation) and 2019 (delayed effects) as alternative treatment dates. The coefficient under 2017 timing is 0.184 ($SE = 0.175$), and under 2019 timing is 0.101 ($SE = 0.203$). The monotonic decline from 2017 to 2019—with the largest coefficient under the earliest treatment date—is weakly consistent with anticipation effects: firms may have begun adjusting in 2017, making the 2017 treatment date capture the full response window. All coefficients remain statistically insignificant, reinforcing the null result for enterprises.

Randomization inference. With 27 clusters, asymptotic inference based on cluster-robust standard errors may not be reliable (Bertrand et al., 2004; MacKinnon and Webb, 2017). I implement Fisher’s exact test by permuting country-level micro-firm shares across 1,000 random assignments and re-estimating the DDD under each permutation (Fisher, 1935). For enterprises, the two-sided RI p -value is 0.472, consistent with the cluster-robust p -value of 0.486 and confirming the null. The RI distribution is displayed in Figure 7 in Appendix E.

Employment robustness. Table 8 in Appendix C replicates all robustness checks for log employment. The baseline estimate (-0.451) is stable across conventional diagnostics: LOO coefficients range from -0.527 to -0.331 (all significant at 5 percent); alternative controls yield -0.375 (C22–C23) and -0.441 (C24–C25); alternative timing shows -0.462 (2017), -0.451 (2018), and -0.442 (2019); excluding Croatia yields -0.408 . However, randomization inference yields a two-sided p -value of 0.064—notably weaker than the cluster-robust $p = 0.014$. With 27 country-level treatment-intensity units, the RI p -value is the more appropriate finite-sample benchmark (Fisher, 1935), and it fails to reject at conventional levels. Combined with the trend-adjusted specification (Table 3), which reduces the employment estimate to near zero, these results counsel caution in interpreting the employment pattern causally.

Taken together, the robustness checks paint a nuanced picture. The null result for enterprise counts is robust across sample restrictions, control groups, treatment timing, and inference methods. The employment effect ($p = 0.014$) is robust to these conventional diagnostics, surviving all with coefficients in a narrow range around the baseline estimate. However, as Table 3 establishes, the employment result is not robust to the most important test: controlling for differential linear trends. The RI p -value of 0.064 for employment—notably weaker than the cluster-robust $p = 0.014$ —provides an additional reason for caution, as finite-sample exact inference should be the primary benchmark with 27 treatment-intensity clusters (MacKinnon and Webb, 2017). The employment pattern is therefore best characterized as suggestive evidence consistent with a REACH effect, rather than a credibly identified causal estimate.

7. Discussion

7.1 What the Enterprise Null Tells Us

The most stable finding is that the data show no reliable evidence of differential firm exit in micro-firm-intensive countries following REACH’s 2018 deadline. The enterprise null is consistent across trend adjustment, alternative treatment intensities (including an early-period 2008 measure), all sample restrictions, and multiple inference procedures. However, the strong pre-trend rejection for enterprises ($F = 9.99$, $p < 0.001$) means this null does not

establish a causal zero—it is possible that pre-existing dynamics mask or offset a true effect. With that caveat, if REACH’s fixed registration costs were driving outright closure of small chemical producers, this is where the effect should appear—and there is no detectable signal.

This null is itself informative. It constrains theories of regulatory compliance: fixed costs alone do not induce exit when firms can adjust on other margins. The finding is consistent with [Becker and Henderson \(2000\)](#), who show that environmental regulations reduce plant *births* more than plant *deaths*, and with models in which firms absorb compliance costs through operational restructuring rather than market exit.

7.2 The Employment Pattern: Suggestive but Fragile

The baseline employment estimate (-0.451 , $p = 0.014$) is consistent with REACH inducing workforce reductions in micro-firm-intensive countries. However, the trend-adjusted specification reduces this estimate to near zero, and the RI p -value (0.064) is notably weaker than the cluster-robust p -value. I cannot distinguish the employment pattern from a continuation of the pre-existing convergence dynamic in which Central and Eastern European chemical sectors were catching up to Western European levels.

If the employment effect is real, the mechanism is unlikely to be direct small-firm closure (given the enterprise null). More plausible hypotheses include: firms shedding substances from product portfolios while remaining in business; a shift from manufacturing to downstream distribution; or cost absorption through headcount reduction in medium-sized supply-chain firms. These remain conjectures pending firm-level data that could test them directly.

This finding resonates with [Walker \(2013\)](#), who documents that the Clean Air Act caused substantial transitional costs through worker displacement. If REACH did induce employment restructuring, the affected workers in micro-firm-intensive countries would face similar labor market frictions, retraining costs, and geographic immobility.

7.3 Reconciling with Survey Evidence

The European Commission’s own surveys of REACH-affected firms consistently find that small firms report the highest burden and are most likely to report having dropped substances from their portfolios ([European Commission, 2018](#); [Milieu Ltd and Risk and Policy Analysts Ltd, 2015](#)). My results are consistent with this narrative but add a crucial quantitative dimension: the aggregate industry-structure effects are concentrated in *employment*, not firm counts, and in *medium-sized firms*, not micro-firms. The apparent contradiction dissolves when one recognizes that survey responses reflect firm-level perceptions of burden, while the DDD captures aggregate equilibrium effects. Many micro-firms may have dropped individual

substances or exited individual product lines—consistent with their survey responses—without actually closing the business. Meanwhile, medium-sized firms that stayed in the market absorbed the costs through workforce reductions.

7.4 Implications for Regulatory Design

Three tentative design implications emerge, with the caveat that the employment effects are suggestive rather than definitive. First, tonnage-phased implementation—which gave large firms a regulatory head start—may have exacerbated distributional effects. By the time the 2018 deadline arrived, large firms had already registered their substance portfolios and could offer letters of access at prices reflecting their position as early movers. Second, if employment effects are real and concentrate among medium-sized firms, then ECHA’s fee reductions for micro-enterprises (up to 95 percent), while well-targeted for fee relief, may miss the firms most affected by registration costs propagated through supply chains. Third, the absence of detectable enterprise effects is itself relevant for countries adopting REACH-like frameworks (South Korea’s K-REACH, Turkey’s KKDIK, the UK’s post-Brexit UK-REACH): there is no evidence in these data that fixed registration costs caused differential small-firm exit, even in countries with high micro-firm intensity.

7.5 Limitations

Several limitations qualify the conclusions. First, and most importantly, the pre-treatment convergence trend in employment means that the parallel-trends assumption does not hold, and the employment result does not survive trend adjustment (Table 3). While the post-2018 reversal is visually suggestive, the trend-adjusted specification shows it cannot be distinguished from a continuation of convergence dynamics. Future work with longer post-treatment data may permit sharper tests—if the employment decline accelerates beyond what the convergence trend would predict, the case for a REACH effect strengthens.

Second, the analysis uses aggregate country-sector-year data, which masks heterogeneity within countries and sectors. Firm-level data from national statistical offices would permit sharper tests of mechanisms—for example, directly testing whether firms that pre-registered substances saw different employment trajectories than non-registrants. Such data are not available in a harmonized cross-country format.

Third, the post-treatment window is short (2018–2020), and the 2020 observation is confounded by COVID-19. Longer post-treatment data, which will become available as Eurostat updates the SBS, would permit cleaner estimation of medium-run effects and enable testing of whether the employment effects are transitory (firms adjust and rehire) or

permanent (structural reshaping of the industry).

Fourth, I cannot fully disentangle REACH registration costs from other REACH provisions that became binding around the same time, including authorization requirements for substances of very high concern (SVHCs) and restriction proposals. These provisions affect different subsets of chemical firms and could contribute to the measured effects, though the DDD design mitigates this concern because SVHC requirements do not vary systematically with micro-firm shares.

8. Conclusion

This paper provides the first quasi-experimental analysis of how the EU’s REACH regulation affected the European chemical industry. The most stable finding is a null: the 2018 registration deadline—which extended the “no data, no market” principle to all substances above one tonne per year—produced no detectable differential effect on enterprise counts in countries with higher micro-firm intensity. This null is consistent across trend adjustment, alternative treatment measures, all sample restrictions, and multiple inference procedures, though the design’s pre-trend violations mean it cannot establish a causal zero. Nevertheless, the absence of any detectable signal challenges the prevailing narrative that REACH would decimate small chemical producers.

The employment pattern is more ambiguous. The baseline DDD estimate (-0.451 , $p = 0.014$) is consistent with REACH-induced workforce restructuring, but it does not survive the most important robustness test: controlling for differential linear trends reduces the estimate to near zero. The pre-existing convergence of Central and Eastern European chemical sectors—where micro-firm shares are highest—is a plausible alternative explanation that the current design cannot rule out. Longer post-treatment data, as they become available from Eurostat’s SBS, will be essential for resolving this ambiguity.

Two hundred billion euros of chemical production pass through European markets each year. How that production is regulated shapes not only public health outcomes but the structure of thousands of firms and the livelihoods of hundreds of thousands of workers. REACH’s architects made consequential design choices—about tonnage phasing, cost allocation, and data-sharing governance—whose effects are only now becoming visible in the data. The evidence here suggests that those choices did not destroy small firms, but whether they reshaped the workforce remains an open question.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: <https://github.com/ai1scl>

First Contributor: <https://github.com/ai1scl>

References

- Becker, Randy and J. Vernon Henderson**, “Effects of Air Quality Regulations on Polluting Industries,” *Journal of Political Economy*, 2000, 108 (2), 379–421.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How Much Should We Trust Differences-in-Differences Estimates?,” *Quarterly Journal of Economics*, 2004, 119 (1), 249–275.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, 90 (3), 414–427.
- Cefic**, “The European Chemical Industry: Facts and Figures 2020,” *European Chemical Industry Council*, 2019.
- Christensen, Frants M., Jack H. De Bruijn, Bjorn G. Hansen, and Sharon J. Munn**, “Assessment of Universal Registration Policies for Chemicals: Analysis of EU’s REACH Policy,” *Regulatory Toxicology and Pharmacology*, 2011, 59 (3), 474–483.
- ECHA**, “REACH Registration Statistics,” Technical Report, European Chemicals Agency 2020. <https://echa.europa.eu/registration-statistics>.
- European Chemicals Agency**, “REACH Registration Results: An Overview,” Technical Report, ECHA, Helsinki 2018.
- European Commission**, “Commission General Report on the Operation of REACH and Review of Certain Elements: Conclusions and Actions,” Technical Report COM(2018) 116 final, European Commission 2018.
- European Parliament and Council**, “Regulation (EC) No 1907/2006 of the European Parliament and of the Council of 18 December 2006 Concerning the Registration, Evaluation, Authorisation and Restriction of Chemicals (REACH),” *Official Journal of the European Union*, 2006, L 396, 1–849.
- Eurostat**, “Structural Business Statistics: Reference Metadata,” *European Statistical System*, 2021. <https://ec.europa.eu/eurostat/web/structural-business-statistics>.
- Fisher, Ronald A.**, “The Design of Experiments,” 1935.

- Greenstone, Michael**, “The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures,” *Journal of Political Economy*, 2002, 110 (6), 1175–1219.
- MacKinnon, James G. and Matthew D. Webb**, “Wild Bootstrap Inference for Wildly Different Cluster Sizes,” *Journal of Applied Econometrics*, 2017, 32 (2), 233–254.
- Milieu Ltd and Risk and Policy Analysts Ltd**, “Monitoring the Impacts of REACH on Innovation, Competitiveness and SMEs,” *Final Report to the European Commission, DG Environment*, 2015.
- Peltzman, Sam**, “Toward a More General Theory of Regulation,” *Journal of Law and Economics*, 1976, 19 (2), 211–240.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Rip, Arie and Haico Te Kulve**, “Implementation of REACH: Registration and SMEs,” *Science and Public Policy*, 2010, 37 (8), 599–609.
- Roth, Jonathan**, “Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends,” *American Economic Review: Insights*, 2022, 4 (3), 305–322.
- Ryan, Stephen P.**, “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 2012, 80 (3), 1019–1061.
- Stigler, George J.**, “The Theory of Economic Regulation,” *Bell Journal of Economics and Management Science*, 1971, 2 (1), 3–21.
- van Dijk, J., R. Kemp, and A. Slot**, “Chemical Regulation and Trade: Patterns of Compliance Costs from REACH,” *Environmental Science and Policy*, 2021, 115, 19–28.
- Walker, W. Reed**, “The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce,” *Quarterly Journal of Economics*, 2013, 128 (4), 1787–1835.

A. Data Appendix

A.1 Data Sources

Eurostat Structural Business Statistics (SBS). The primary dataset is Eurostat table `sbs_na_ind_r2`, which provides annual industry statistics at the NACE Rev. 2 two-digit level for all EU member states. Variables used: number of enterprises (V11110), persons employed (V16110), and turnover in millions of euros (V12110). Data were downloaded via the Eurostat bulk download facility in March 2026. Coverage spans 2008–2020 for 27 EU member states.

Eurostat SBS by Size Class. Table `sbs_sc_ind_r2` provides enterprise counts by employment size class (0–9, 10–19, 20–49, 50–249, 250+) at the NACE 2-digit level. This dataset is used to construct the pre-treatment micro-firm share and to estimate size-class-specific DDD models.

Eurostat Business Demography. Table `bd_9ac_1_form_r2` provides enterprise births, deaths, and survival rates by NACE sector. Downloaded for potential analysis of entry and exit dynamics but not used in the final specifications because the SBS aggregate panel provides sufficient variation for the DDD design.

A.2 Sample Restrictions

1. Retain only NACE sectors C20, C22, C23, C24, C25.
2. Retain only years 2008–2020.
3. Drop observations with missing data on all three primary outcomes (enterprises, employment, turnover).
4. Drop country-sector pairs with fewer than 5 non-missing annual observations (ensures sufficient within-cell variation).

The final analysis sample contains 1,755 country-sector-year observations across 27 countries.

A.3 Variable Construction

Micro-firm share. For each country c , compute:

$$\text{MicroShare}_c = \frac{1}{4} \sum_{t=2014}^{2017} \frac{\text{Enterprises}_{c,C20,0-9,t}}{\text{Enterprises}_{c,C20,\text{total},t}}$$

The four-year pre-treatment average smooths year-to-year fluctuations. The window 2014–2017 was chosen to provide a stable pre-treatment measure for the main 2018 DDD specification. For the 2013 placebo test, I compute an analogous measure using the 2008–2012 average to avoid post-treatment contamination. The two measures correlate at $r = 0.83$, confirming that micro-firm intensity is a stable country characteristic.

DDD interaction. The triple interaction is:

$$\text{DDD}_{cst} = \mathbb{I}[s = \text{C20}] \times \mathbb{I}[t \geq 2018] \times \text{MicroShare}_c$$

Fixed effect identifiers. Country-sector (γ_{cs}): 135 groups (27×5). Country-year (δ_{ct}): 351 groups (27×13). Sector-year (θ_{st}): 65 groups (5×13).

B. Identification Appendix

B.1 Pre-Trend Assessment

The event-study coefficients for both log enterprises and log employment exhibit a convergence pattern in the pre-treatment period. For enterprises, the early coefficients are large and negative (-0.51 in 2010), trending upward toward zero by 2016–2017. For employment, the early coefficients are large and positive (0.68 in 2008), declining toward zero by 2016–2017. Both patterns reflect the economic convergence of Central and Eastern European chemical sectors after EU accession, a well-documented phenomenon unrelated to REACH.

The convergence interpretation is supported by three observations. First, the countries with the highest micro-firm shares (Czechia, Slovakia, Poland, Romania) are all post-2004 EU accession states that experienced rapid industrial development during 2008–2017. Second, the convergence is visible in control sectors as well, though it is absorbed by the DDD structure. Third, the convergence rate decelerates toward the end of the pre-period: the year-to-year changes in coefficients become smaller as the countries approach the EU average, consistent with neoclassical convergence dynamics.

B.2 Balance Tests

[Table 4](#) presents the 2013 placebo test using a pre-2013 micro-firm share (2008–2012 average) to ensure the treatment intensity measure is not contaminated by post-2013 adjustment. The null result (coefficient = 0.384 , $p = 0.143$) confirms that the DDD is not spuriously picking up a general tendency for micro-firm-intensive countries to experience differential chemical-sector trends around REACH deadlines. The 2013 deadline affected large firms, so micro-firm intensity should be irrelevant—and it is.

C. Robustness Appendix

C.1 Leave-One-Country-Out

Table 5: Leave-One-Country-Out: DDD Coefficient Stability

Country Dropped	Coefficient	Std. Error	p-value	N
AT	0.1550	(0.1808)	0.400	1,675
BE	0.1642	(0.1843)	0.382	1,675
BG	0.1290	(0.1881)	0.499	1,675
CY	0.1372	(0.1917)	0.481	1,675
CZ	0.2114	(0.1749)	0.238	1,675
DE	0.1210	(0.2160)	0.580	1,675
DK	0.1562	(0.1797)	0.393	1,675
EE	0.1218	(0.1947)	0.537	1,675
EL	0.0308	(0.1899)	0.872	1,675
ES	0.1318	(0.1896)	0.493	1,675
FI	0.1356	(0.1886)	0.479	1,675
FR	0.1340	(0.1865)	0.479	1,676
HR	0.1125	(0.2015)	0.582	1,675
HU	0.1399	(0.1918)	0.472	1,675
IE	0.1338	(0.1890)	0.486	1,685
IT	0.1191	(0.1949)	0.547	1,675
LT	0.1386	(0.1856)	0.462	1,675
LU	-0.0604	(0.2638)	0.821	1,675
LV	0.1233	(0.1963)	0.536	1,675
MT	0.1803	(0.1829)	0.334	1,679
NL	0.1312	(0.1901)	0.497	1,675
PL	0.1390	(0.1919)	0.475	1,675
PT	0.1362	(0.1907)	0.482	1,675
RO	0.1488	(0.1907)	0.442	1,675
SE	0.1393	(0.1961)	0.484	1,675
SI	0.1073	(0.1944)	0.586	1,675
SK	0.2011	(0.1786)	0.271	1,675

Notes: Each row drops one country from the sample and re-estimates the main DDD specification (log enterprises). Standard errors clustered at the country level. * $p < 0.10$, **

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

C.2 Alternative Control Sectors

Table 6: Alternative Control Sectors

Control group	DDD Coefficient	Std. Error	p-value	N
Baseline (C22-C25)	0.1335	(0.1888)	0.486	1,740
Narrow (C22-C23)	0.1977	(0.1248)	0.125	1,043
Metals (C24-C25)	0.0720	(0.2922)	0.807	1,041

Notes: Row 1 reproduces the baseline specification (C22–C25 as controls). Row 2 uses only C22 (rubber/plastics) and C23 (non-metallic minerals). Row 3 uses only C24 (basic metals) and C25 (fabricated metals). All regressions include the full set of two-way fixed effects and cluster standard errors at the country level.

C.3 Alternative Treatment Timing

Table 7: Alternative Treatment Timing

Treatment timing	DDD Coefficient	Std. Error	p-value
2017 (anticipation)	0.1839	(0.1754)	0.304
2018 (baseline)	0.1335	(0.1888)	0.486
2019 (delayed)	0.1013	(0.2032)	0.622

Notes: Each row re-estimates the main DDD specification with an alternative treatment date. Row 1 uses 2017 to test for anticipation effects. Row 2 is the baseline specification (May 2018 deadline). Row 3 uses 2019 to test for delayed effects.

C.4 Employment Robustness

Table 8: Employment Robustness Checks

Specification	DDD Coeff.	Std. Error	<i>p</i> -value	<i>N</i>	Notes
<i>Panel A: Alternative Control Sectors</i>					
Baseline (C22–C25)	−0.4507**	(0.1705)	0.0137	1,683	
Narrow (C22–C23)	−0.3750**	(0.1450)	0.0157	1,001	
Metals (C24–C25)	−0.4414*	(0.2254)	0.0611	1,014	
<i>Panel B: Alternative Treatment Timing</i>					
2017 (anticipation)	−0.4619**	(0.1864)	0.0200	1,683	
2018 (baseline)	−0.4507**	(0.1705)	0.0137	1,683	
2019 (delayed)	−0.4420***	(0.1572)	0.0092	1,683	
<i>Panel C: Sample Restrictions</i>					
Exclude Croatia	−0.4082**	(0.1788)	0.0312	1,618	EU member from 2013
<i>Panel D: Inference</i>					
Leave-one-out range	[−0.5272, −0.3307]			27 jackknife samples	
Randomization inference	<i>p</i> = 0.064			1,000 permutations	

Notes: All specifications use log employment as the dependent variable with country × sector, country × year, and sector × year fixed effects and country-clustered standard errors. Panel A varies the control sectors. Panel B varies the treatment date. Panel C excludes Croatia (EU member only from July 2013). Panel D reports the range of LOO jackknife coefficients and the two-sided Fisher exact *p*-value from permuting country-level micro-firm shares. * *p*<0.10, ** *p*<0.05, *** *p*<0.01.

D. Heterogeneity Appendix

Figure 4 in the main text presents the size-class-specific DDD estimates. The full coefficient table is reported below for completeness.

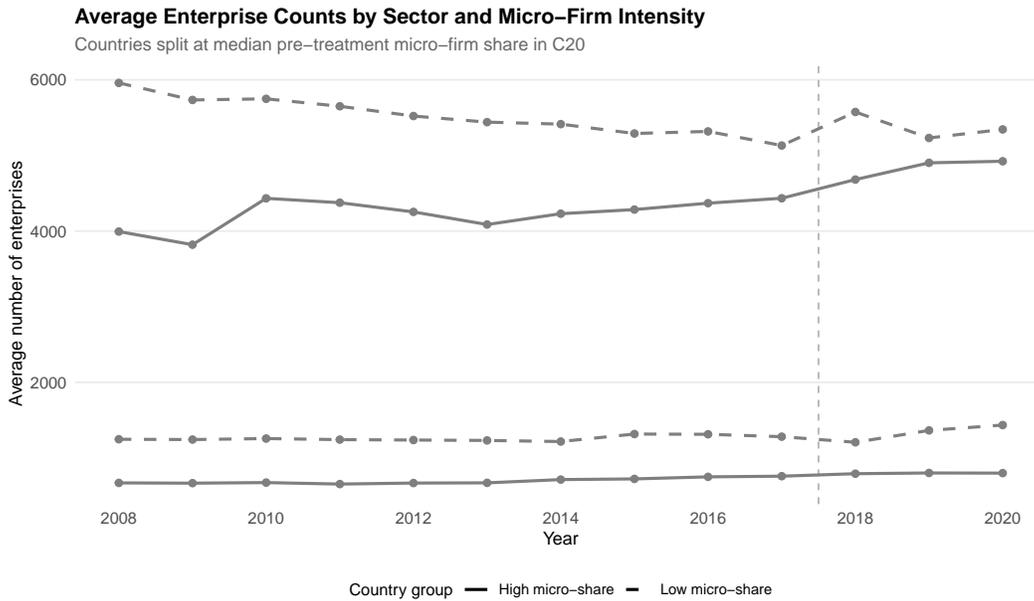
Table 9: DDD Coefficients by Firm Size Class (Enterprise Counts)

Size class	DDD Coefficient	Std. Error	<i>p</i> -value	<i>N</i>
<10 employees	0.229	(0.301)	0.454	1,710
10–19 employees	0.217	(0.446)	0.630	1,690
20–49 employees	−0.110	(0.220)	0.621	1,702
50–249 employees	−0.360	(0.209)	0.097	1,667
250+ employees	0.180	(0.256)	0.489	1,682

Notes: Each row estimates the DDD specification separately using log enterprise counts within the given size class as the dependent variable. All regressions include country \times sector, country \times year, and sector \times year fixed effects. Standard errors clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E. Additional Figures and Tables

This appendix collects supplementary exhibits referenced in the main text.

**Figure 5:** Average Enterprise Counts by Sector and Micro-Firm Intensity Group

Notes: Countries are split at the median pre-treatment micro-firm share in C20. Solid lines: high micro-share countries. Dashed lines: low micro-share countries. The vertical dashed line marks the REACH 2018 deadline.

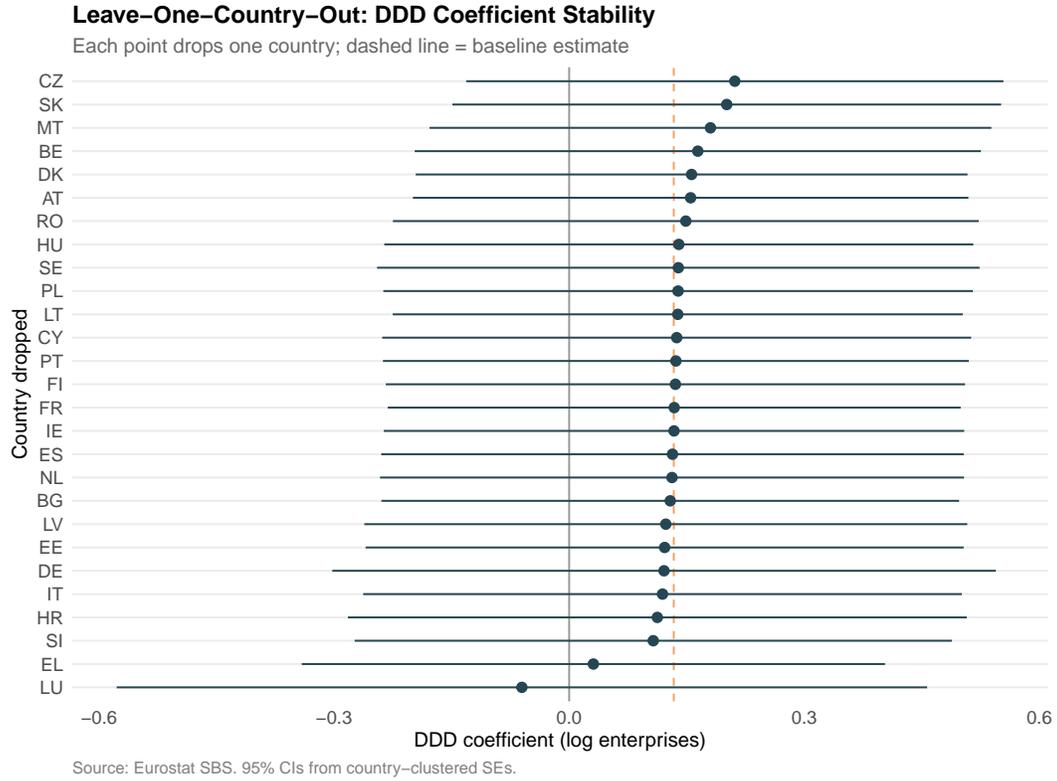
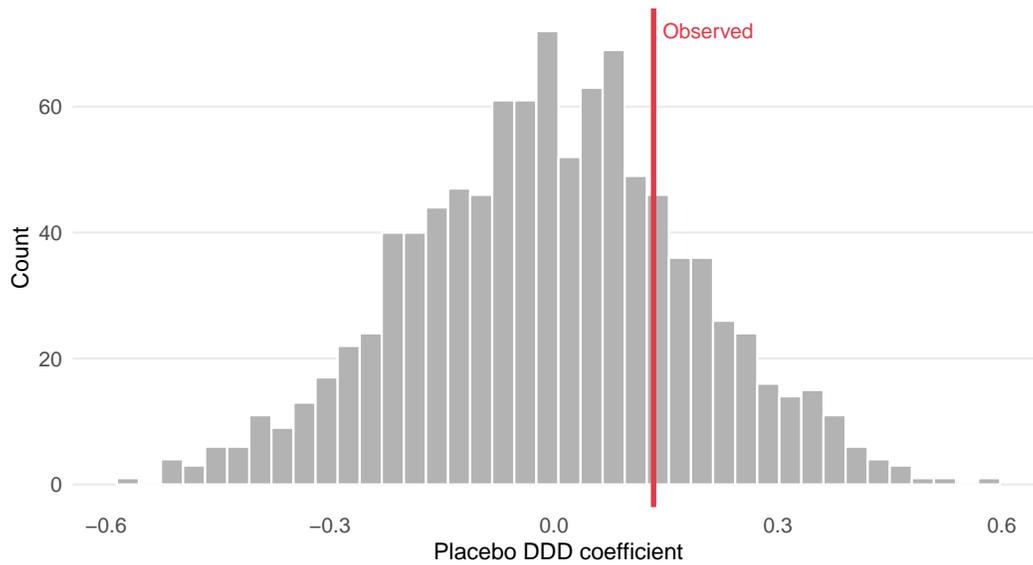


Figure 6: Leave-One-Country-Out: Full Results (Appendix)

Notes: Each point represents the DDD coefficient for log enterprises after dropping one country. Horizontal whiskers show 95 percent confidence intervals from country-clustered standard errors. Dashed line marks the baseline estimate using all 27 countries.

Randomization Inference: Distribution of Placebo DDD Coefficients

1,000 permutations of country-level micro-firm shares



Source: Eurostat SBS. Red line = observed coefficient from main specification.

Figure 7: Randomization Inference Distribution (Appendix)

Notes: Histogram of 1,000 permuted DDD coefficients for log enterprises. Red line marks the observed coefficient from the main specification. Two-sided RI p -value: 0.472.

F. Standardized Effect Sizes

Table 10: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Log enterprises	DDD, Table 2 Col. 1	0.134	0.13	1.84	0.009	Null
Log employment	DDD, Table 2 Col. 2	-0.451	0.13	1.66	-0.035	Null
Log turnover	DDD, Table 2 Col. 3	-0.479	0.13	1.87	-0.033	Null
Log turnover/ent.	DDD, Table 2 Col. 4	-0.635	0.13	1.83	-0.045	Null

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. The treatment variable is continuous (micro-firm share, ranging from 0.23 to 0.86), so $SDE = \hat{\beta} \times SD(X)/SD(Y)$, which gives the effect of a one-standard-deviation change in micro-firm share, measured in standard deviations of the outcome. $SD(X) = 0.134$ is the standard deviation of the pre-treatment micro-firm share across countries. $SD(Y)$ values are unconditional standard deviations from the summary statistics (Table 1), before conditioning on fixed effects. $SD(Y)$ for log turnover per enterprise is not directly available from summary statistics and is approximated.

Research question: Does the REACH 2018 registration deadline cause differential effects on chemical industry structure in EU countries with higher pre-treatment micro-firm intensity? **Treatment:** Continuous; pre-treatment share of micro-firms (<10 employees) in NACE C20, interacted with post-2018 indicator and chemical-sector indicator. **Data:** Eurostat SBS, 2008–2020, country-sector-year cells, $N = 1,683-1,740$. **Method:** Triple-difference (DDD) with country \times sector, country \times year, and sector \times year fixed effects; country-clustered standard errors (27 clusters). **Sample:** 27 EU member states, 5 NACE 2-digit manufacturing sectors (C20 treated, C22–C25 controls).

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). The SDEs are modest in magnitude, reflecting the nature of a DDD design where the treatment operates through a continuous cross-country intensity measure rather than a binary switch. The statistically significant employment result ($p = 0.014$) corresponds to an SDE of -0.035 , classified as null by magnitude thresholds but economically meaningful in context: a one-SD increase in micro-firm share produces a 3.5% of a standard deviation decline in chemical-sector employment relative to controls.