CSEF

Centre for Studies in Economics and Finance

WORKING PAPER NO. 784

Does Common Carbon Pricing Reorient Trade? Partner Selection and the EU ETS

Angelo De Santis, Carla Guerriero, Antonia Pacelli, and Carmine Russo

June 2026

University of Naples Federico II

University of Salerno

Bocconi University, Milan

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS AND STATISTICS – UNIVERSITY OF NAPLES FEDERICO II
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 – e-mail: csef@unina.it
ISSN: 2240-9696

CSEF
Centre for Studies in Economics and Finance

WORKING PAPER NO. 784

Does Common Carbon Pricing Reorient Trade? Partner Selection and the EU ETS

Angelo De Santis[*], Carla Guerriero[†], Antonia Pacelli[‡], and Carmine Russo[§]

Abstract
We study whether participation in the EU Emissions Trading System (EU ETS) reshapes the selection of trading partners, rather than merely raising production costs. A common carbon price works through two opposing channels: a cost channel that discourages trade with regulated countries, and a regulatory-alignment channel that lowers policy uncertainty between co-regulated partners. Combining a gravity framework with a difference-in-differences design over 1995–2020, we find that bilateral imports between EU ETS members rise by about 11 to 33% relative to non-participants (33% in a staggered Callaway–Sant'Anna design) while exports show no net effect. The import premium grows as the carbon price rises across policy phases and, in product-level data, is mostly concentrated in ETS-covered sectors, pointing to regulatory alignment rather than general integration. The results show that common carbon pricing can strengthen trade integration among insiders without necessarily reorienting trade away from unregulated partners.

JEL Classification: C23, F14, F18, Q54, Q58.

Keywords: climate club; regulatory alignment; trading-partner selection; carbon pricing; gravity model; emission trading system; carbon leakage.

Acknowledgments: We thank Saverio Simonelli, IAERE, SIEP and EAERE conference participants for insightful comments and suggestions. The research is funded by the European Union - Next Generation EU, in the framework of the GRINS - Growing Resilient, Inclusive and Sustainable Project (GRINSPE00000018 – CUP E63C22002140007). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, and the European Union cannot be held responsible for them.

[*] University of Naples Federico II.

[†] University of Naples Federico II and CSEF. E-mail: carla.guerriero@unina.it.

[‡] Toulouse School of Economics and INRAE, University of Toulouse Capitole. E-mail: antonia.pacelli@tse-fr.eu *(Corresponding author)*

[§] Vienna University of Economics and Business. E-mail: carmine.russo@wu.ac.at *(Corresponding author)*

1 Introduction

The EU ETS, launched in 2005, is the world's largest multinational carbon market and currently the second-largest emissions trading scheme globally, following China's national ETS. It covers approximately 45% of the European Union's total greenhouse-gas emissions and is a cornerstone of EU climate policy (EU Parliament, 2003). By placing an explicit price on carbon, it creates cost-effective incentives for reducing emissions (Dechezleprêtre et al., 2023; Colmer et al., 2025). Because the scheme currently applies to only 30 countries (the 27 EU member states plus Iceland, Norway and Liechtenstein), it also creates an asymmetric regulatory environment in global markets, in which some trading relationships fall entirely within a common carbon-pricing regime while others straddle its boundary. The trade consequences of this asymmetry have been debated primarily through the lens of competitiveness and carbon leakage, that is, whether carbon pricing pushes trade away from regulated economies. We ask a different question: whether common carbon pricing changes with whom regulated countries trade.

This paper studies whether participation in the EU ETS is associated with a relative reorientation of trade toward partners subject to the same carbon-pricing regime. We interpret such a pattern through the lens of regulatory alignment: when importer and exporter operate under a common carbon-pricing architecture, future cost exposure and compliance conditions become more predictable, and the risk of future divergence, such as border-adjustment measures applied to outsiders, is reduced. In the presence of sunk costs of building and maintaining trade relationships, lower policy uncertainty can raise trade with similarly regulated partners (Handley and Limão, 2015, 2017). In this sense the EU ETS may operate as a de facto climate club (Nordhaus, 2015; Ernst et al., 2023): a common carbon-pricing regime can generate an "insider" trade premium through regulatory convergence and lower intra-club frictions, even without explicit penalties on non-members. Unlike the original climate-club framework, our analysis does not rely on explicit sanctions or trade penalties against non-participants. Instead, we examine whether a common carbon-pricing regime itself can foster trade integration among insiders.

This logic implies that carbon pricing affects bilateral trade through two channels that operate in opposite directions. On the one hand, carbon pricing raises the marginal cost of regulated firms and may discourage trade involving regulated economies, the standard carbon-leakage force. On the other hand, a common regulatory framework may reduce uncertainty and implicit frictions between participating countries, encouraging trade among co-regulated partners. Alignment is naturally relevant on the sourcing (import) margin, where firms select suppliers under common regulatory conditions; the cost channel is naturally relevant on the competitiveness (export) margin, where it directly affects regulated exporters. The net effect is therefore theoretically ambiguous and must be determined empirically. We develop this framework in Section 3.

We embed a gravity model in a difference-in-differences design. The sample covers the EU ETS countries as origins over 1995–2020 together with their main trading partners, retaining for each origin-year the top 10% of partners by bilateral trade on each margin.[1] The baseline analysis restricts attention to the 25 countries that joined the scheme at its 2005 launch and uses 2000, the year of the policy's official announcement, as the reference period, following Colmer et al. (2025), on the grounds that forward-looking firms may begin adjusting once the policy is announced.[2] We complement the baseline with event studies, estimated both on the early-adopter sample and on the full set of 30 ETS countries, and with a staggered difference-in-differences estimator (Callaway and Sant'Anna, 2021) that assigns countries to cohorts by the year their participation was

[1]For each EU ETS origin country and year, we rank all trading partners by bilateral import (export) volume and retain the top 10% within that country-year distribution, separately on each margin. The filter is applied origin by origin, so each country contributes its own most relevant partners (about 20 per country-year on average), preserving country-specific trade structures.

[2]The 25 first-phase countries are Austria, Belgium, Cyprus, the Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, Poland, Portugal, the Slovak Republic, Slovenia, Spain, Sweden and the United Kingdom. The United Kingdom is treated as an EU ETS member throughout because it participated until the end of 2020, the last year of our sample, and is dropped in a robustness check. Analogous results using 2005 as the reference year are reported in Appendix C.

announced and so handles the heterogeneous entry timing of later entrants.[3]

Our main result is that EU ETS participation raises bilateral imports between regulated countries relative to non-participants by about 11 to 33%, while the effect on exports is not statistically different from zero. As we show, this aggregate export null is not an absence of effect but the netting-out of opposing sector-specific forces, and the underlying mechanism (regulatory alignment within the sectors the scheme covers) operates on both margins. Event-study estimates show that the import effect emerges after the policy and strengthens over time, becoming most pronounced in the more stringent Phase III; the export response remains flat throughout. The staggered estimator, which gives more weight to long horizons and to later cohorts, delivers an aggregate effect of about 33% on imports and essentially zero on exports.

A natural concern is that this import premium reflects general trade integration among EU members (the single market and the 2004 enlargement, which coincide with the post-announcement period) rather than carbon regulation as such. Our most direct evidence against this reading is a sector-level analysis built from Base pour l'Analyse du Commerce International (BACI) product-level data. Replicating the design separately for sectors covered by the EU ETS and for structurally non-covered downstream sectors, we find the import premium concentrated in a covered sector (basic iron and steel) and absent in an otherwise comparable downstream sector (fabricated metal products) that shares the same buyers, trade geography and price signals but operates no regulated installations. The same contrast appears on exports with opposite signs: a positive intra-ETS premium in the regulated good and a negative effect in the non-covered downstream sector, so that the two offset in the aggregate and explain the export null. Because general integration would lift both sectors on both margins, this contrast points to a mechanism specific to the activities the ETS actually regulates, namely regulatory alignment, rather than a broad bilateral proximity effect.

[3]Later entrants are assigned to cohorts by the announcement of their participation, which for these countries is close to their accession or EEA-entry year: Bulgaria and Romania (2006), Norway and Iceland (2007), and Croatia (2012). Liechtenstein is excluded throughout because of missing trade data.

The findings are robust to a battery of checks. First, they are not driven by common-currency dynamics: the intra-ETS import premium survives when the sample is restricted to pairs in which at least one partner has not adopted the euro.[4] Second, they are not an artefact of the non-independence of mirrored ETS pairs: clustering at the undirected-dyad level rather than the directed pair leaves the estimates and their precision virtually unchanged. Third, they are not an artefact of the top-10% partner filter: removing it and re-estimating on all bilateral partners yields, if anything, a slightly larger and more precisely estimated import coefficient, so that the filter delivers a more conservative estimate (Section 6.3). Fourth, the results are stable when we additionally control for the share of origin-country emissions covered by the ETS (Appendix E). The results are also robust to estimating the gravity equation in multiplicative form by Poisson pseudo-maximum likelihood (PPML) (Santos Silva and Tenreyro, 2006), which accommodates zero trade flows and the heteroskedasticity of log-linearized trade data. This specification provides the most conservative coefficient (11%) on imports compared to the log-OLS and staggered estimates (about 20% and 33%). This result is moreover stable whether or not the top-partner filter is applied; all specifications agree on a positive and significant import premium and on a null on exports.

These results indicate that participation in the EU ETS strengthens trade integration among member countries on the import margin, with no sign of an aggregate reorientation toward unregulated partners. By harmonizing carbon-cost exposure and reducing policy uncertainty across regulated economies, the scheme appears to lower implicit trade frictions within the regulated group, consistent with a de facto climate club. These effects are likely reinforced by firms' anticipation of future border-adjustment measures such as the EU's Carbon Border Adjustment Mechanism (CBAM), which raise the expected cost of sourcing from non-regulated regions (Chen, 2023). Our results should not be read as a comprehensive test of carbon leakage: the relative shift toward ETS partners does not by itself rule out leakage at a more disaggregated level or in embodied-emission terms.

[4]See Appendix D.

Rather, it shows that, at the aggregate bilateral level, trade does not reorient toward unregulated partners as the dominant response to carbon pricing.

Our paper contributes to three strands of the literature. First, to the literature on trade and the EU ETS, which has overwhelmingly asked whether carbon pricing pushes trade away from regulated economies through leakage: Antoci et al. (2021) model leakage as a firm response to regulatory asymmetries, Wang and Kuusi (2024) find with a sector-level gravity model that EU import carbon intensity rises while export intensity falls, and direct tests detect little aggregate leakage (Naegele and Zaklan, 2019; Branger et al., 2016), in contrast to the leakage that Aichele and Felbermayr (2015) document for Kyoto signatories outside the ETS.[5] We reverse the question: rather than asking whether trade flows away from regulated countries, we ask whether it flows toward co-regulated partners, and we document a robust insider import premium. We also complement Hintermann and Ludwig (2023), who apply a gravity framework to the EU ETS allowance market and uncover a strong home-country bias rooted in pre-existing goods-trade patterns; we turn instead to goods trade itself and ask how common regulation reshapes it.

Second, we contribute to the gravity literature on how policy regimes reshape bilateral trade. Methodologically, our panel difference-in-differences with country-pair and year fixed effects follows the approach Baier and Bergstrand (2007) developed to identify the trade effects of free trade agreements, applied here to a common regulatory regime rather than a tariff arrangement. Our question is with whom regulated countries trade, and it connects to the literature on the selection of trading partners and the extensive margin of trade (Helpman et al., 2008); in the main analysis, we document how a common regulatory regime reshapes sourcing along the intensive margin of bilateral volume within the set of top trading partners, leaving the extensive-margin response (whether the ETS alters the total number of partners) to future work. Substantively, a large quantitative-trade literature studies the trade and welfare effects of carbon tariffs and border adjustments (Larch and Wanner, 2017; Ernst et al., 2023); we complement this largely general-equilibrium

[5]For a review of leakage measurement, see Fowlie and Reguant (2018).

literature, which emphasises the cost and leakage channels and the incentives to join a carbon-pricing club, by documenting a reduced-form alignment channel that reshapes the geography of trade even in the absence of explicit border instruments. Third, we contribute to the literature on policy uncertainty and trade (Handley, 2014; Handley and Limão, 2015, 2017), providing evidence that a common and stable regulatory regime raises trade among insiders, consistent with the option-value and sunk-cost mechanisms at the heart of that work.

Third, our paper relates to the large literature evaluating the EU ETS. This literature documents substantial emission reductions from the scheme (Dechezleprêtre et al., 2023; Colmer et al., 2025; Biancalani et al., 2024) without detectable losses in output, employment or profitability (Colmer et al., 2025; Bremer and Sommer, 2025), and points to low-carbon innovation and cleaner entry among regulated firms (Calel and Dechezleprêtre, 2016; Chen, 2023; Pacelli, 2025). This evidence speaks to the competitiveness debate framed by the competing "pollution haven" (De Beule et al., 2022; Dechezleprêtre and Sato, 2017) and "Porter" (Porter and Linde, 1995; Ambec et al., 2013) hypotheses, from which our partner-selection question departs.[6]

The remainder of the paper is organized as follows. Section 2 describes the institutional setting. Section 3 develops the two-channel framework. Section 4 presents the data and the difference-in-differences strategy. Section 5 reports the main results, including the sector-level evidence on the mechanism. Section 6 presents the robustness analysis. Section 7 concludes.

[6] A related macro literature studies the aggregate effects of carbon pricing, generally finding limited effects of carbon taxes on GDP, employment and inflation (Metcalf, 2019; Bernard and Kichian, 2021; Metcalf and Stock, 2023; Konradt and Weder, 2021), while ETS-based pricing identified from policy surprises can generate larger macroeconomic and distributional effects (Känzig, 2023; Mangiante, 2024). Gravity and DiD models have also been applied to ETS and other settings, such as aviation demand (Oesingmann, 2022; Falk and Hagsten, 2019).

2 Policy context

The EU ETS is the European Union's main instrument for decarbonizing the economy. It is a market-based policy tool, complex and continuously evolving as a result of learning by doing and responses to unforeseen events. It currently covers 30 countries across Europe (the 27 EU Member States plus Iceland, Liechtenstein, and Norway) and places limits on carbon dioxide (CO_2), nitrous oxide (N_2O), and perfluorocarbons (PFCs) emissions from more than 11,000 energy-intensive installations (EU Parliament, 2003). Owing to its central role in European climate policy and its relevance in the global climate governance framework, the academic literature on the EU ETS has expanded rapidly over the past two decades.

The policy was announced in 2000, the EU Directive was adopted in 2003, and the scheme became operational in 2005. Since then, the EU ETS has developed through four distinct trading phases, each characterized by increasing regulatory stringency and institutional refinement (European Commission, 2020). Phase I, spanning from 2005 to 2007, functioned primarily as a pilot period to establish the institutional and market infrastructure necessary for carbon trading and to test price formation mechanisms. The system initially included the EU-25 member states, with Bulgaria and Romania joining in 2007 upon EU accession. Nearly all emission allowances were allocated for free under the grandfathering principle to protect industrial competitiveness and to fight carbon leakage concerns (Woerdman et al., 2008). However, the pilot phase resulted in a substantial allowance surplus and a collapse in carbon prices to nearly zero by the end of the phase (Narassimhan et al., 2018). Verified emissions were considerably below allocated allowances, underscoring the limited environmental impact of the initial phase but establishing the operational framework for future phases. Covered sectors included large combustion plants, oil refineries, iron and steel production, cement clinker, glass, lime, bricks, ceramics, pulp, paper, and board.

Phase II, 2008-2012, coincided with the Kyoto Protocol's first commitment period and introduced more stringent reduction requirements. Nitrous oxide (N_2O) was added to the list of regulated gases, and the penalty for non-compliance increased from €40 to €100 per tonne of CO_2-equivalent (European Commission, 2009). The European Economic Area countries: Norway, Iceland, and Liechtenstein joined the ETS during this phase, and the aviation sector was included in 2012. While free allocation remained the dominant distribution method, monitoring, reporting, and verification standards were strengthened (Ellerman and Buchner, 2008). Carbon prices initially traded around €20 per tonne but fell to around €10 in later years, reflecting persistent over-allocation.

Phase III, which lasted until 2020, marked a major reform of the ETS architecture, harmonizing the cap and allocation rules at the EU level and replacing national allocation plans with a single EU-wide cap. The default allocation method shifted from free allocation to auctioning, significantly reducing the potential for windfall profits. Coverage expanded to include additional industrial activities such as aluminium production, petrochemicals, ammonia, and various acid production processes, as well as carbon capture, transport, and geological storage.

Phase III also built on the institutional and market reforms of the earlier phases and laid the ground for the alignment of the ETS with the EU's commitments under the Paris Agreement. The annual linear cap-reduction factor was raised, and the Market Stability Reserve[7] began operating in 2019. The share of free allocation continued to decline, with auctioning becoming the predominant distribution method. The evolution of the EU ETS demonstrates a steady progression toward greater environmental effectiveness, market integration, and policy credibility. While early phases were characterized by over-allocation and weak price signals, later phases have introduced harmonized caps, auctioning, and market stabilization mechanisms, transforming the ETS into a central and increasingly effective component of the EU's climate strategy. Appendix B reports the historical evolution of the EU allowance price across the four phases.

[7] The Market Stability Reserve is a mechanism within the EU Emissions Trading System designed to automatically adjust the supply of emission allowances in the carbon market in order to maintain market stability and prevent large imbalances between supply and demand.

Two features of this institutional history are central to our empirical strategy. First, the scheme was announced in 2000, before it became operational in 2005, which is why we use 2000 as the reference year. Second, the sharp rise in stringency across phases, from near-zero allowance prices in the pilot phase to a tightening cap and markedly higher prices in Phase III, provides the time variation we exploit in the phase-by-phase analysis. If carbon pricing affects trade through the cost channel, its bite should grow as the phases progress, whereas the alignment channel is present as soon as the common regime exists; the phase profile therefore helps separate the two.

3 Theoretical framework: regulatory alignment and bilateral trade

We use the structural gravity model as a reduced-form framework for bilateral trade and extend it to let carbon pricing shape trade frictions. Appendix A reviews the theoretical foundations of the gravity equation and the empirical literature on the distance elasticity of trade. In standard gravity, bilateral trade between countries i and j rises with their economic size and falls with bilateral trade costs (Anderson and Van Wincoop, 2003):

$$X_{ijt} = \exp\left(\alpha_{it} + \delta_{jt} - \tau_{ijt}\right), \tag{1}$$

where α_{it} and δ_{jt} collect importer-time and exporter-time factors, including multilateral resistance, and τ_{ijt} denotes bilateral trade frictions. Our interpretation works entirely through τ_{ijt}: we ask how common carbon pricing reshapes the frictions between two countries.

3.1 A two-channel framework for carbon pricing and trade

A common carbon-pricing institution such as the EU ETS can affect these frictions through two channels of opposite sign. The first works through marginal production costs. By pricing emissions, the ETS raises the marginal cost of regulated firms; standard trade theory then predicts that higher costs reduce the competitiveness of regulated exporters and make imports from non-regulated, lower-cost suppliers more attractive. Operating alone, this cost channel diverts trade away from regulated economies, reflecting the standard carbon-leakage mechanism emphasized in both empirical and quantitative trade models (Aichele and Felbermayr, 2015; Ernst et al., 2023).

The second works through regulatory alignment. Common participation in the ETS reduces asymmetries in carbon-cost exposure across partners: when both i and j price carbon under the same architecture, future compliance conditions are more predictable, regulatory environments more comparable, and the risk of future divergence, such as border-adjustment measures applied to outsiders, is lower. In the presence of sunk costs of establishing and maintaining trade relationships (Handley and Limão, 2015, 2017), this convergence lowers implicit frictions among co-regulated partners and, operating alone, encourages trade within the regulated group.

3.2 Reduced-form representation

We capture both channels by decomposing bilateral frictions into a time-invariant geographic component and a time-varying regulatory component:

$$\tau_{ijt} = \tau_{ij}^{geo} + \tau_{ijt}^{reg}, \tag{2}$$

where τ_{ij}^{geo} collects persistent bilateral barriers such as distance and language, and τ_{ijt}^{reg} the regulatory frictions linked to environmental policy. We write the latter as

$$\tau_{ijt}^{reg} = \underbrace{\phi \cdot \mathrm{MC}_{ijt}}_{\text{Channel 1: cost}} + \underbrace{\psi \cdot \mathrm{AsymReg}_{ijt}}_{\text{Channel 2: alignment}}, \tag{3}$$

where MC_{ijt} summarises the marginal-cost impact of carbon pricing on the pair (larger when at least one partner is regulated and carbon prices are high) and $\mathrm{AsymReg}_{ijt}$ the difference in carbon-cost exposure between i and j (smaller when both belong to the ETS). The two components enter intra-ETS trade with opposite signs: the cost channel ($\phi > 0$) raises frictions involving regulated countries, while the alignment channel ($\psi > 0$, with AsymReg_{ijt} lower for ETS–ETS pairs) lowers them.

3.3 Interpretation and testable implications

Our empirical strategy estimates a difference-in-differences coefficient comparing bilateral trade between ETS and non-ETS partners before and after the introduction of the EU ETS. This coefficient measures the net effect of the two channels described above. The cost channel tends to reduce trade involving regulated countries, whereas the regulatory-alignment channel tends to increase trade between co-regulated partners. The overall effect is therefore theoretically ambiguous. A positive coefficient indicates that the alignment channel dominates the cost channel on average, while a negative coefficient would suggest the opposite. Determining which channel prevails is ultimately an empirical question, and the reduced-form coefficient is informative precisely because it measures the net balance rather than presupposing it.

Although the framework does not deliver a prediction for the sign of the average treatment effect, it generates three empirical implications that help distinguish the regulatory-alignment mechanism from alternative explanations. First, timing. Any effect attributable to the EU ETS should emerge only after the announcement of the policy and not before. Prior to the announcement, trade between future ETS and non-ETS partners should evolve similarly. We assess this prediction through event-study estimates, which also provide a test of the parallel-trends assumption. Second, EU ETS phase dynamics. The economic relevance of carbon pricing changes substantially across ETS phases. During the pilot phase, allowance prices were close to zero, implying a limited role for carbon costs. Both channels become stronger as carbon pricing becomes more stringent. The cost wedge rises mechanically with the allowance price, while the benefits of regulatory alignment increase because future compliance conditions become more predictable and regulatory asymmetries more costly (e.g. the introduction of the CBAM). The evolution of the estimated trade premium across phases therefore provides information on how the balance between the cost and regulatory-alignment channels changes as carbon pricing becomes more consequential. Since both channels scale with carbon-price stringency, the evolution of the estimated premium across phases provides information about their relative importance. A growing premium implies that the regulatory-alignment channel expands at least as fast as the cost channel. Third, sectoral concentration. If the estimated trade premium reflects regulatory alignment under the EU ETS, it should be concentrated in sectors that are covered by the scheme, where firms face common carbon-pricing rules. Comparable downstream sectors that share similar buyers, trade geography, and market conditions but are not directly regulated should exhibit a weaker response. By contrast, explanations based on general European integration would predict similar effects across both covered and non-covered activities. We exploit this distinction in Section 5.4 to assess whether the observed trade premium is linked to common carbon pricing, discriminating a regulatory-alignment mechanism from a broad proximity or integration effect.

The two-channel framework provides guidance on how the effects may differ across import and export margins. The regulatory-alignment channel is expected to operate most directly on the sourcing (import) side, where firms, choosing among suppliers, favour partners operating under a common carbon-pricing regime. The cost channel, by contrast, affects the competitiveness of regulated producers and therefore affects exports. We highlight that in our directional design the exporter's own carbon cost does not vary with the partner's regulatory status, so the net export effect is a priori indeterminate and is better read in light of the sector-level evidence, where opposing effects in covered and non-covered sectors can offset in the aggregate (Section 5.4).

4 Data and methodology

This section describes the construction of the dataset and the empirical strategy used to assess the impact of the EU ETS on trading-partner selection. The empirical analysis is anchored to the gravity model introduced in Section 3.

4.1 Sample construction

Our panel covers the EU ETS countries as origins over 1995–2020 together with their main trading partners. Liechtenstein is excluded throughout for missing trade data, leaving 30 origin countries (29 + United Kingdom); the baseline further restricts attention to the 25 first-phase adopters. For each origin and year we rank all partners by bilateral trade volume and keep the top 10% on each margin, separately for imports and exports. This origin-specific rule lets each country contribute its own most relevant partners (about 20 per country-year), preserving country-specific trade structures and limiting noise from marginal relationships; Section 6.3 shows that dropping the filter and using all bilateral partners leaves the results unchanged. On average, this procedure retains approximately 20 partners per country-year, both inside and outside the EU ETS, yielding meaningful variation across the treated–control comparison.

The bilateral panel is directional. The unit of observation is a country-pair-year (i, j, t), where i is an EU ETS origin country and j is a trading partner; for each origin country we observe both its imports from and its exports to each partner. Because origins are restricted to EU ETS countries, the panel differs between ETS–ETS and ETS–non-ETS pairs. For an ETS–ETS pair such as Italy–Germany the relationship appears in both directions, so that the import flow on one side closely mirrors the export flow on the other (the correlation between $\log(\text{import}_{ij})$ and $\log(\text{export}_{ji})$ across mirrored ETS–ETS observations exceeds 0.97); we account for this non-independence by clustering on the undirected dyad in Section 6.2. For an ETS–non-ETS pair such as Italy–China only the direction with the ETS country as origin is observed. We estimate the import and export equations separately because they capture distinct margins, sourcing on the import side and competitiveness on the export side (Section 3), and because the ETS–non-ETS pairs that identify the effect carry different information on the two margins. Export flows are also intrinsically noisier than imports in this setting, which is relevant when interpreting the export estimates.

4.2 Data sources and variables

Bilateral imports and exports are measured in thousands of US dollars and come from the World Bank World Integrated Trade Solution (WITS) database.[8] Data on GDP (constant 2015 US dollars, thousands) are obtained from the World Bank. Bilateral distance is the population-weighted harmonic mean distance between the two countries' largest cities, time-variant and in kilometres, from Centre d'Études Prospectives et d'Informations Internationales (CEPII). Environmental taxes, from the OECD, capture the stringency of environmental policy beyond the ETS itself. Further time-invariant bilateral covariates (common language, contiguity, common legal origin, and the religious-proximity index are available from CEPII but are absorbed by the country-pair fixed effects and so are not included as regressors. From Our World in Data and the European Environment Agency we additionally construct, for each origin and year, the share of national emissions covered by the ETS, used as a control in a robustness specification (Appendix E).

Table 1 reports summary statistics for the full sample.

Trade flows display substantial heterogeneity across country pairs. While average log imports and exports are 8.1 and 9.3, respectively, the observed values span more than 23 log points, ranging from very small trade relationships to some of the largest bilateral flows in the sample. This considerable variation provides substantial identifying power for the gravity analysis. Economic size, as captured by GDP at origin and destination, is more tightly distributed, although the range still reflects the presence of both small and large economies in the sample, ensuring meaningful variation in market size. Environ-

[8]Public dataset available at this link.

	Full			
	Mean	SD	Min	Max
Log(import)	8.113	4.474	-4.605	18.720
Log(export)	9.301	3.559	-4.605	18.771
$Log(GDP_{origin})$	26.538	1.491	22.869	29.278
$Log(GDP_{partner})$	25.104	2.240	17.209	30.911
Log(env. taxes_{origin})	8.447	1.687	1.698	11.315
Log(env. $\text{taxes}_{partner}$)	6.618	2.535	-4.564	11.952
Log(Distance)	8.367	0.899	4.007	9.882
Observations	145557			

Table 1: Descriptive statistics, full sample.

mental taxation also exhibits non-negligible dispersion: while average levels are relatively similar across origin and partner countries, the range points to substantial cross-country differences in environmental policy stringency. Geographical distance is the most dispersed variable, with values ranging from 55 km (Austria–Slovak Republic) to nearly 19,000 km for the most distant partners.

4.3 Empirical strategy

We use a difference-in-differences (DiD) design that compares bilateral trade before and after the 2000 announcement of the EU ETS, distinguishing partners that joined the scheme (treated) from those that remained outside (control). As mentioned, we take 2000 rather than the 2005 start as the reference year, following Colmer et al. (2025), because forward-looking firms may begin adjusting at announcement.[9] The baseline restricts origins to the 25 first-phase adopters.

Our main specification is

$$\log(Y_{ijt}) = \beta_1 \cdot \text{did}_{ijt} + \beta_2 \cdot \log(\text{GDP}_{jt}) + \beta_3 \cdot \log(\text{GDP}_{it}) + \beta_4 \cdot \text{Dist}_{ij} + \beta_5 \cdot \text{EnvTax}_{jt} + \beta_6 \cdot \text{EnvTax}_{it} + \gamma_{ij} + \delta_t + \varepsilon_{ijt},$$

[9] Appendix C reports analogous results using 2005 as the reference year.

where Y_{ijt} is either the import or the export flow between country i and partner j in year t, expressed in thousands of US dollars. The key independent variable, did_{ijt}, is an indicator equal to one if the trading partner j belongs to the EU ETS and the year is after the policy announcement. We control for the economic size of both countries through log(GDP), for geographic harmonic distance (Dist), and for the environmental taxes applied in both countries (EnvTax). Country-pair fixed effects γ_{ij} absorb time-invariant bilateral determinants of trade, such as common language, contiguity, common legal origin and the religious-proximity index of Disdier and Head (2008); year fixed effects δ_t absorb common annual shocks. Bilateral distance is retained as a regressor because the population-weighted harmonic measure varies over time, so it is not collinear with the pair effects; its coefficient is identified only from this within-pair variation and is not the object of interest. We estimate the import and export equations separately. In a further specification we also control for the share of emissions covered by the EU ETS in the origin country, accounting for the fact that the policy bites only on the portion of national emissions that the scheme actually regulates.[10]

Treated and control pairs differ on observables (Table 2): treated pairs trade more and are markedly closer, with somewhat different partner-GDP and tax levels. These imbalances motivate the controls, while the within-pair design absorbs their level differences.

	Baseline		T1		T0			
	Mean	SD	Mean	SD	Mean	SD	T0-T1	t
Log(imports)	8.259	4.489	12.939	2.583	7.606	4.307	-5.333***	(-215.734)
Log(exports)	9.467	3.545	12.949	2.648	8.946	3.363	-4.003***	(-167.015)
Log(GDP_{origin})	26.674	1.505	26.493	1.547	26.699	1.497	0.206***	(15.444)
Log($\text{GDP}_{partner}$)	25.060	2.261	26.552	1.534	24.797	2.267	-1.756***	(-119.739)
Log(env tax_{origin})	8.623	1.704	8.401	1.787	8.654	1.689	0.253***	(16.427)
Log(env $\text{tax}_{partener}$)	6.588	2.548	8.438	1.768	5.884	2.445	-2.555***	(-135.330)
log(Distance)	8.370	0.899	7.103	0.658	8.589	0.740	1.486***	(241.462)
Observations	123812		15161		108651		123812	

T1 and T0 represent respectively the treated and control group.

Table 2: Descriptive statistics, baseline sample.

The baseline difference-in-differences estimates are complemented by several addi-

[10] See Appendix E.

tional analyses designed to assess the timing and dynamics of the estimated effect. First, we estimate event-study specifications to trace the evolution of the treatment effect and test the parallel-trends assumption. These are implemented both on the baseline sample of early adopters and on the full set of 30 EU ETS countries. Second, we exploit the institutional evolution of the EU ETS through a phase-by-phase analysis, examining how the estimated trade premium changes as carbon pricing becomes progressively more stringent. Third, we implement the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), which exploits variation in the timing of participation across countries and complements the common-reference-year design used in the baseline specification.[11] Finally, we consider a specification that additionally controls for the share of national emissions covered by the EU ETS in the origin country (Appendix E).

5 Results

This section presents the main results of the analysis. We begin with the baseline difference-in-differences on the early adopters, trace the dynamics with an event study and its phase-by-phase version, extend the event study to the full set of ETS countries, and apply the staggered estimator of Callaway and Sant'Anna (2021) for heterogeneous entry timing. A final subsection turns to sector-level evidence on the mechanism. Throughout, the estimates bear on the three restrictions of Section 3.3: timing and parallel trends (event study), the phase profile (phase by phase), and sectoral analysis (Section 5.4).

5.1 Baseline results: early EU ETS adopters

The baseline focuses on the 25 countries that participated in the EU ETS from its inception in 2005 as origin, with country-pair fixed effects. As in Colmer et al. (2025), flows are benchmarked against the announcement year (2000). Table 3 reports the estimates: EU

[11] Liechtenstein is excluded throughout because of missing trade data. The United Kingdom is treated as an EU ETS member because it participated in the scheme until the end of 2020, the final year of our sample.

ETS membership raises bilateral imports between participating countries by about 20% relative to non-participants,[12] while the export coefficient is small and not statistically different from zero. The positive import effect indicates that, on the sourcing margin, the alignment channel dominates the cost channel; we return to the export null in Section 5.4, where the sector-level evidence shows it masks offsetting effects across covered and non-covered sectors rather than a clean within-relationship cancellation. As expected in a gravity setting, both origin and partner GDP carry large, precisely estimated elasticities. The choice of the top-10% filter is not driving the result: the import effect remains positive and significant as the threshold is relaxed, and under PPML is stable whether or not the filter is applied (Sections 6.3 and 6.4).

The contrast between the import and export coefficients deserves comment, because intra-ETS trade enters both equations. Every intra-ETS import is, by construction, an intra-ETS export recorded with origin and partner reversed, so the treated flows are essentially common to the two equations (the correlation between mirrored intra-ETS imports and exports exceeds 0.97). The two coefficients therefore differ not because intra-ETS imports and exports behave differently, but because each is benchmarked against a different control group: imports from non-ETS partners in one case, exports to non-ETS partners in the other. The difference-in-differences premium is a relative object, so a positive import premium together with a null export premium says that, relative to trade with the rest of the world, ETS countries shifted their sourcing toward co-regulated partners while their sales did not tilt differentially toward them. This is consistent with partner selection operating on the sourcing margin.[13]

Figures 1 and 2 show that treated and control pairs followed similar trajectories before 2000, with no evidence of differential pre-trends. This graphical evidence supports the parallel-trends assumption underlying the difference-in-differences design and complements the baseline estimates reported above. The results suggest that EU ETS

[12] With a log dependent variable and a binary treatment, percentage effects are $\%\Delta Y = (e^{\widehat{\beta}} - 1) \times 100\%$.

[13] As we show next, the within-sector comparison provides cleaner evidence on the mechanism (Section 5.4).

	(1) log_imp	(2) log_exp
did	0.184***	0.0482
	(3.29)	(0.73)
$\log(GDP_{country})$	1.058***	0.527***
	(6.00)	(3.40)
$\log(GDP_{partner})$	0.627***	0.963***
	(3.73)	(6.00)
log(distance)	-11.46**	-6.234
	(-2.72)	(-1.36)
$\log(\text{env tax}_{partner})$	0.480***	0.204**
	(7.65)	(3.14)
$\log(\text{env tax}_{country})$	0.205***	0.489***
	(5.20)	(9.88)
FE year	✓	✓
FE id	✓	✓
Observations	11050	11365

t statistics in parentheses; standard errors clustered at the country-pair level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Baseline results.

participation is associated with a positive reorientation of imports toward countries operating under the same carbon-pricing regime. The findings are also unchanged when we additionally control for the share of origin emissions covered by the ETS (Appendix E). While the baseline estimates establish the existence of this asymmetry, they do not by themselves reveal the underlying mechanism. The next subsections therefore examine the timing and evolution of the effect and then turn to sector-level evidence to assess whether it reflects regulatory alignment rather than broader forms of economic integration.

5.2 Event studies baseline and full sample

The event-study estimates enrich the baseline findings. Figures 3 and 4 plot the dynamic import and export coefficients relative to 2000. On the import side, the coefficients are small and insignificant before 2000, turn positive after the policy, and are largest in the more stringent Phase III. On the export side they remain close to zero throughout with

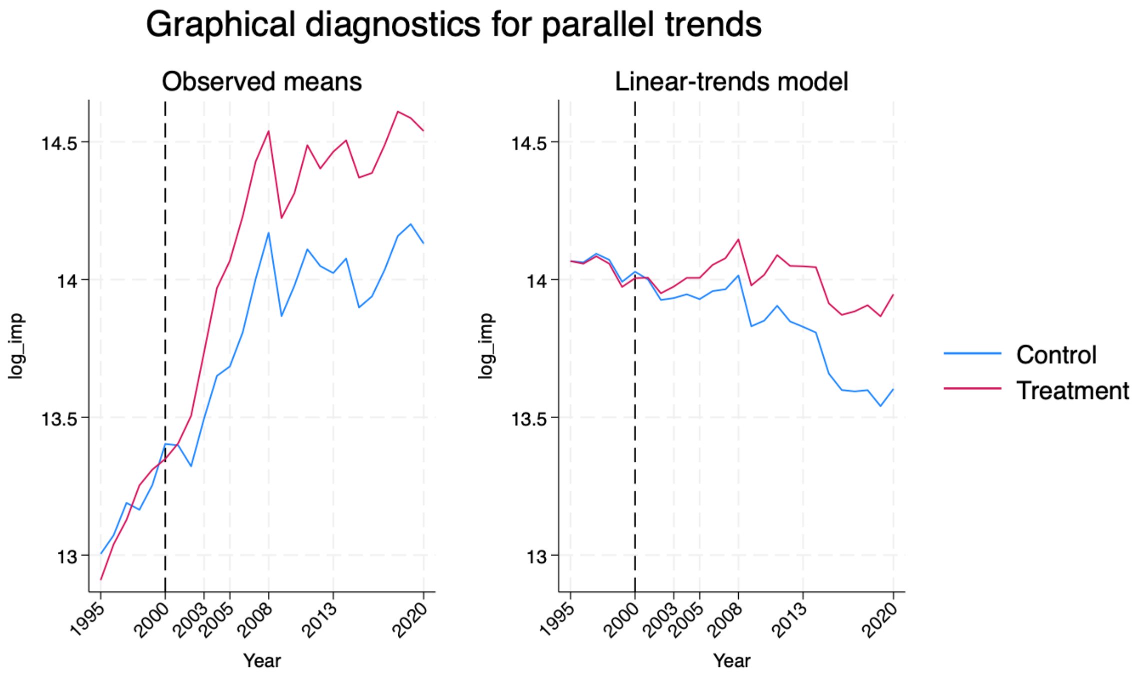

Figure 1: Baseline sample: log import in treated and control group.

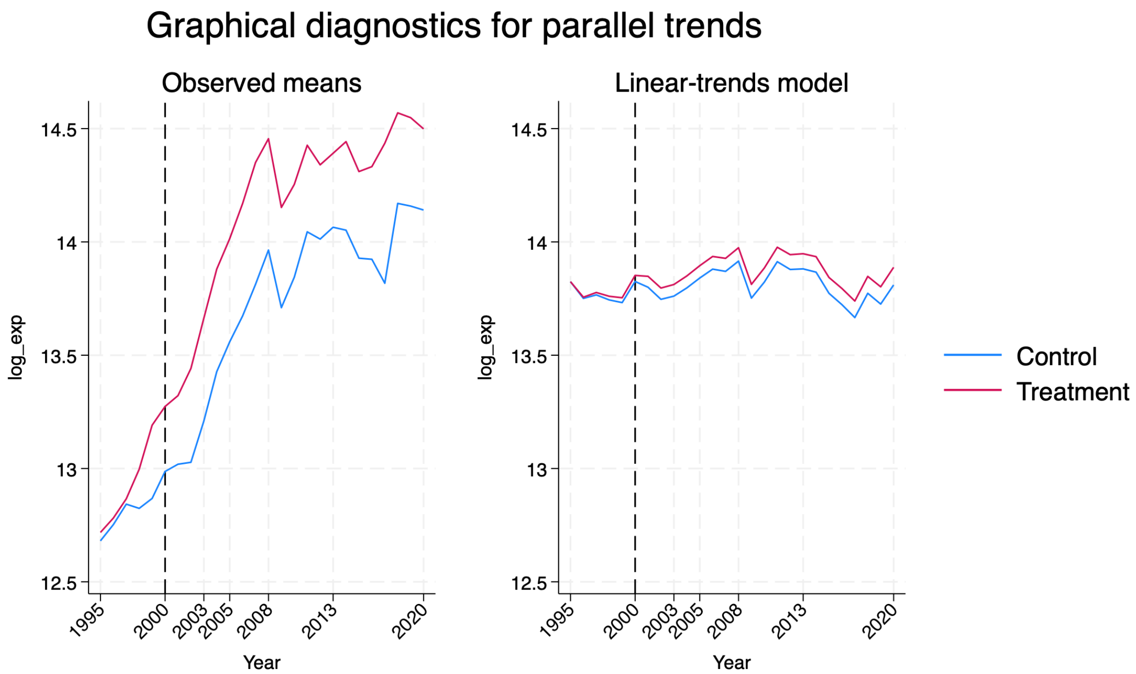

Figure 2: Baseline sample: log export in treated and control group.

no discernible post-treatment pattern. The negligible pre-2000 coefficients provide direct support for parallel trends.

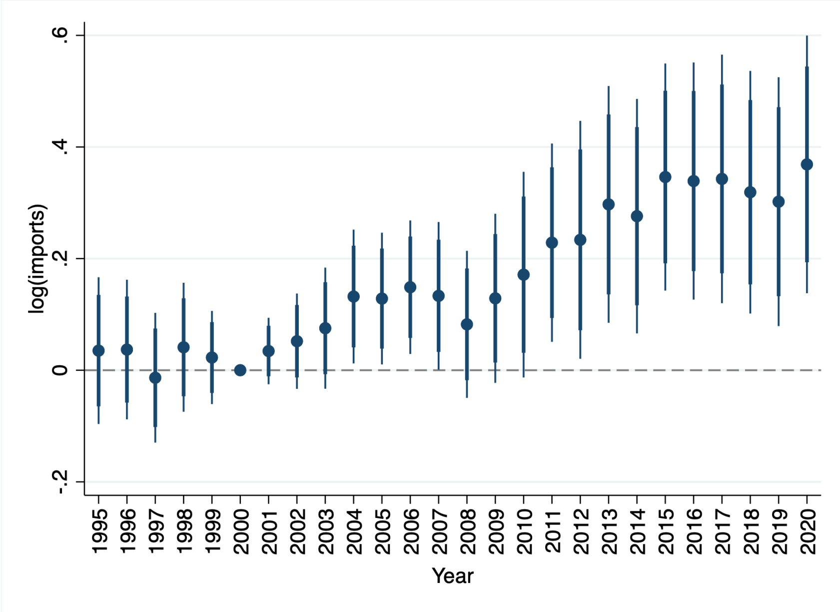

Figure 3: Baseline sample: event study for imports.

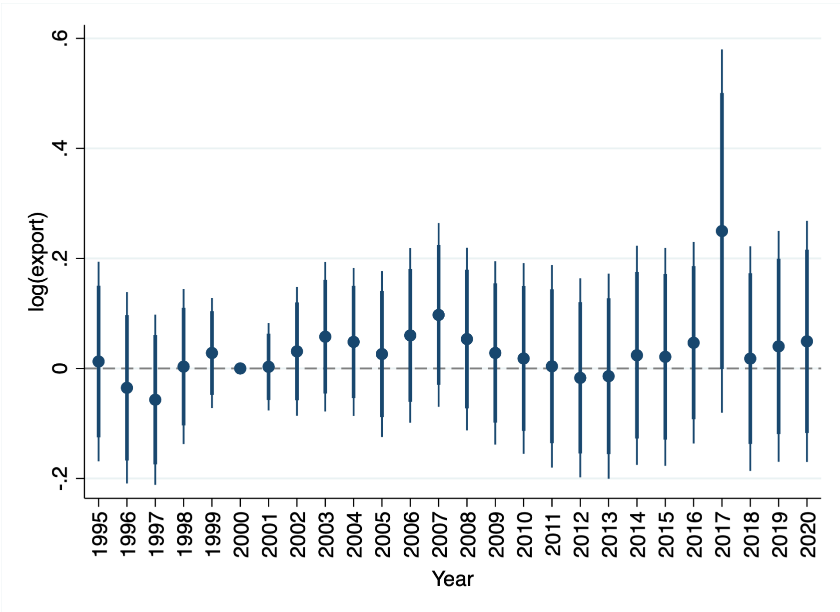

Figure 4: Baseline sample: event study for exports.

5.2.1 Phase-by-phase analysis

To study the effects phase by phase, we estimate the following regression:

$$\log(Y_{ijt}) = \alpha + \sum_{p=1}^{6} \beta_p \cdot T_j \cdot \mathbf{1}\{t \in \text{Phase}_p\} + \pi_1 \cdot \log(\text{GDP}_{jt}) + \pi_2 \cdot \log(\text{GDP}_{it}) + \pi_3 \cdot \text{Dist}_{ij}$$
$$+ \pi_4 \cdot \text{EnvTax}_{jt} + \pi_5 \cdot \text{EnvTax}_{it} + \gamma_{ij} + \delta_t + \varepsilon_{ijt},$$

In line with Colmer et al. (2025), $\{\text{Phase}_p\}_{p=1}^{6}$ is an indicator equal to one if year t belongs to phase p and zero otherwise; interacted with T_j, the resulting phase-specific DiD coefficients β_p measure the trade differential between ETS and non-ETS pairs in each phase relative to the announcement-year reference (2000).[14]

The import response rises monotonically across phases (Figures 5 and 6): it is negligible before and immediately after the announcement, and becomes positive and significant from the early 2000s, growing monotonically to its maximum in Phase III. The premium is already positive and significant in the pilot phase, when carbon prices remained low and the cost implications of the scheme were still limited. This pattern is consistent with the regulatory-alignment channel operating as soon as a common carbon-pricing regime is established, with the effect becoming stronger as carbon pricing grows more stringent in later phases. The export response is flat and insignificant in every phase, and pre-announcement coefficients (1995–1999) are small and statistically indistinguishable from zero for both margins, supporting the parallel-trends assumption. This is precisely the phase profile anticipated in Section 3.3: the net effect grows with stringency, consistent with the alignment channel continuing to outweigh the rising cost wedge over the sample period. As carbon prices rose in Phase II and especially Phase III, the premium widened further, consistent with the alignment channel scaling with stringency and the anticipated cost of sourcing from non-regulated partners.

We replicate the same event-study analysis using the full sample, including all countries that are subject to the EU ETS, regardless of the phase in which they entered the

[14] The phases are defined as: $Phases_1 = \{1995, \ldots, 1999\}$, $Phases_2 = \{2001, \ldots, 2002\}$, $Phases_3 = \{2003, \ldots, 2004\}$, $Phases_4 = \{2005, \ldots, 2007\}$, $Phases_5 = \{2008, \ldots, 2012\}$, $Phases_6 = \{2013, \ldots, 2020\}$.

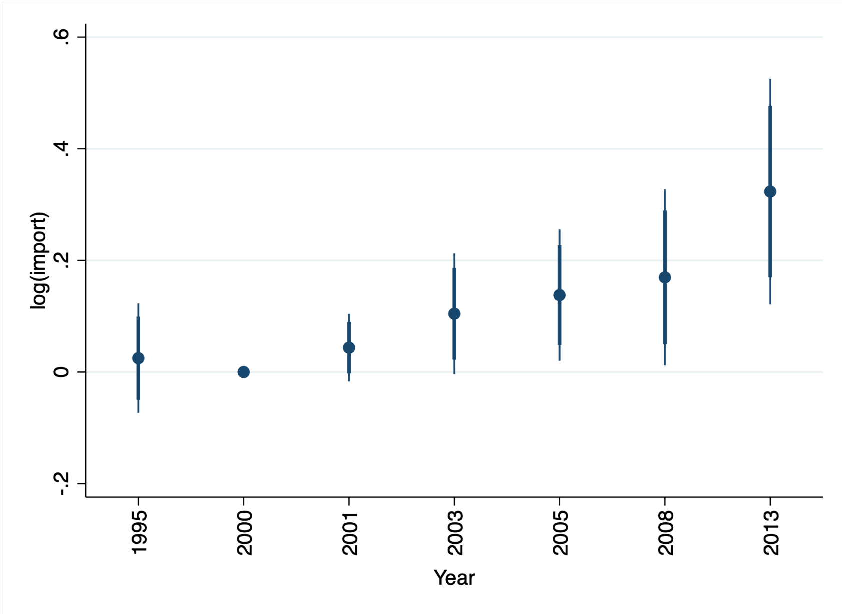

Figure 5: Baseline sample: phase-by-phase event study for imports.

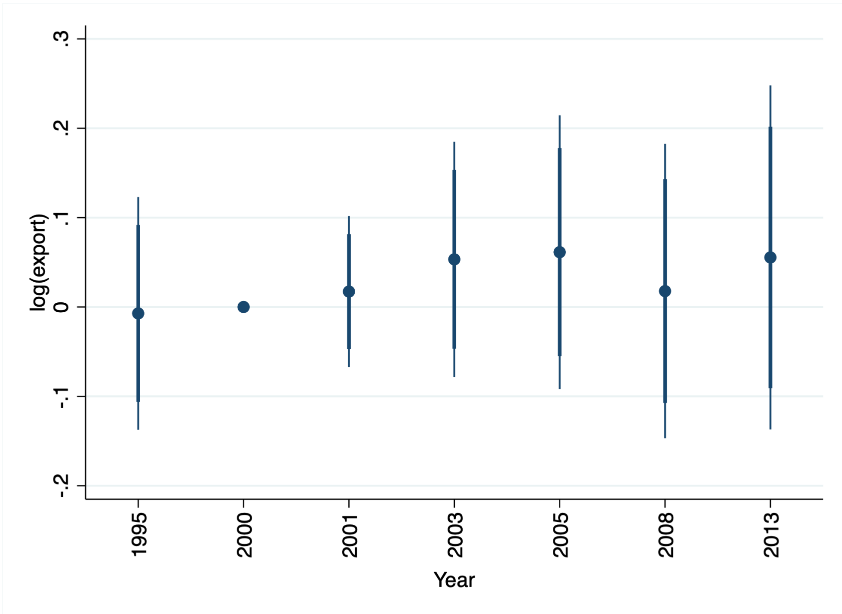

Figure 6: Baseline sample: phase-by-phase event study for exports.

scheme. The resulting estimates, presented in Figures 7 and 8, are qualitatively and quantitatively similar to those obtained in the restricted sample. In particular, the dynamics, significance and import-export contrast are quantitatively similar to the early-adopter sample, indicating that the findings are not driven by first-phase entrants. Because this specification treats all ETS countries as a single group with a common 2000 reference, however, it is coarse for later entrants, whose treatment begins much later; the staggered estimator in the next subsection addresses this directly.

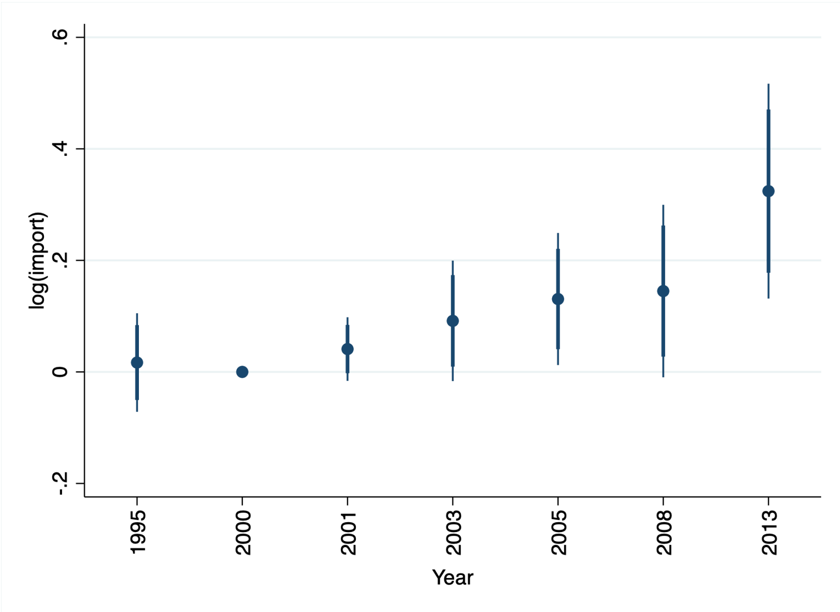

Figure 7: log import event study of EU ETS phases full sample.

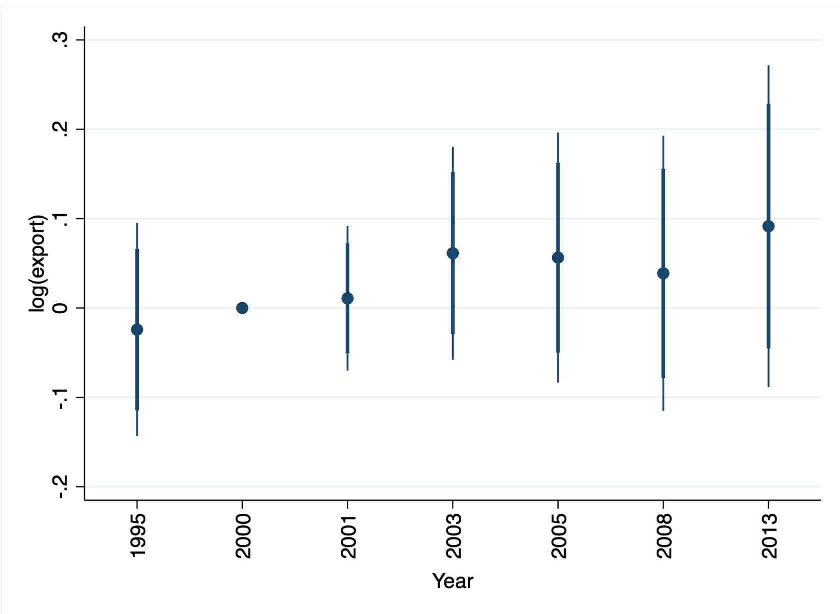

Figure 8: log export event study of EU ETS phases full sample.

5.3 Staggered difference-in-differences

We exploit the heterogeneous timing of entry with the Callaway and Sant'Anna (2021) estimator, assigning countries to cohorts by the announcement of their participation.[15] Figure 9 and 10 report the group-time average treatment effects (ATTs) on trade flows. The full-sample descriptives underlying this exercise are those in Table 4.

	Full		T1		T0			
	Mean	SD	Mean	SD	Mean	SD	T0-T1	t
Log(Import)	8.113	4.474	12.748	2.574	7.444	4.290	-5.304***	(-235.677)
Log(Export)	9.301	3.559	12.701	2.679	8.770	3.382	-3.931***	(-177.683)
Log(GDP origin)	26.538	1.491	26.347	1.529	26.566	1.484	0.219***	(18.177)
Log(GDP partner)	25.104	2.240	26.551	1.535	24.843	2.248	-1.709***	(-127.870)
Log(Env. tax. origin)	8.447	1.687	8.220	1.745	8.480	1.676	0.260***	(18.920)
Log(Env. tax partner)	6.618	2.535	8.438	1.770	5.911	2.433	-2.527***	(-146.800)
log(distance)	8.367	0.899	7.130	0.653	8.586	0.745	1.457***	(261.748)
Observations	145557		18322		127235		145557	

T1 and T0 represent respectively the treated and control group.

Table 4: Descriptive full sample

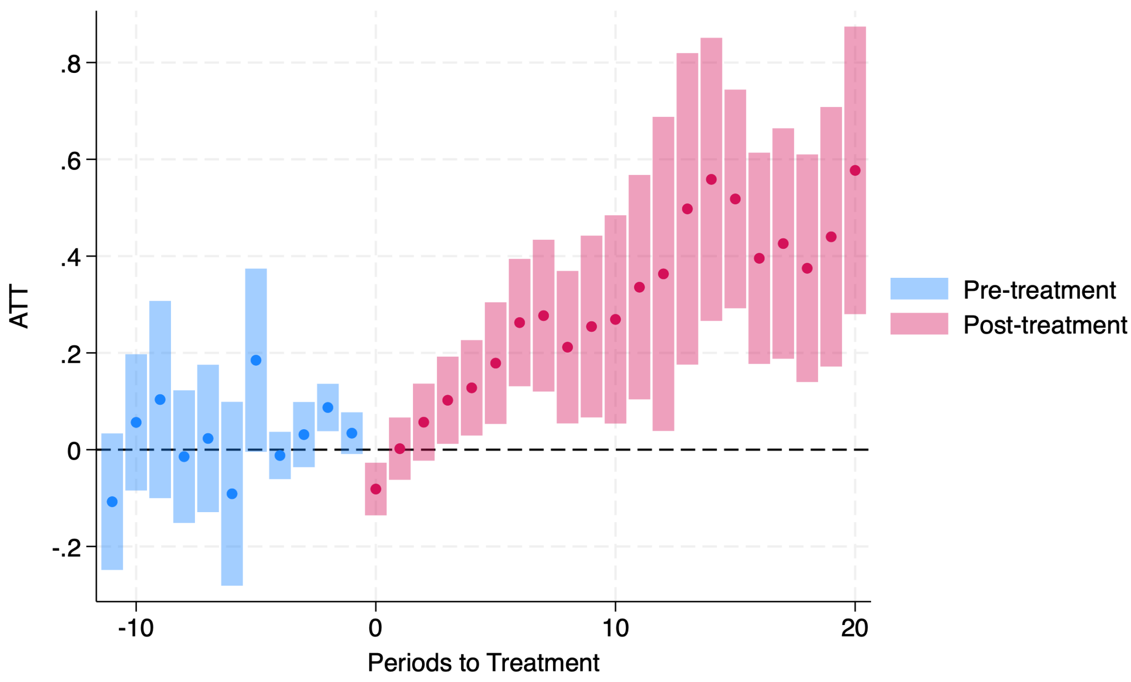

Figure 9: Log of import (US dollars) difference between EU ETS vs non-EU ETS participants.

[15]2000 for first-phase adopters; 2006 for Bulgaria and Romania; 2007 for Norway and Iceland; 2012 for Croatia

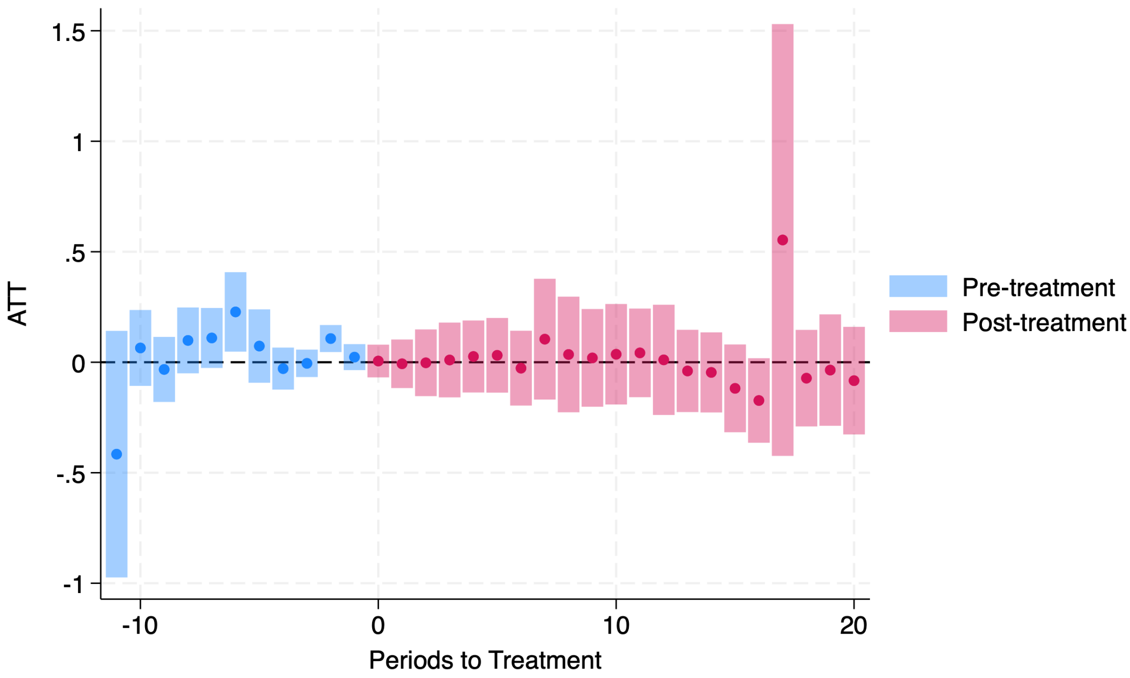

Figure 10: Log of exports (US dollars) difference between EU ETS vs non-EU ETS participants.

Relative to the baseline sample, the full sample displays very similar patterns, confirming that the main differences between treated and control groups are not driven by sample selection. Trade flows remain substantially higher among treated observations, with imports and exports close to one log point larger than in the control group. Both the magnitude and statistical significance of these differences are remarkably stable, indicating that treated pairs are consistently characterized by stronger trade linkages. The comparison in terms of economic size also closely mirrors the baseline. Differences in GDP at origin remain small, although now precisely estimated, while partner-country GDP continues to be significantly lower in the treated group. This suggests that the tendency for treated observations to involve relatively smaller destination markets is a robust feature of the data. A similar stability emerges for environmental taxation. As in the baseline, treated countries exhibit slightly lower environmental taxes at origin and higher levels at destination. The magnitude of these gaps remains modest but statistically significant, reinforcing the presence of systematic policy differences between treated and control units. Geographical distance continues to represent the most pronounced dimension of divergence.

The consistency between the baseline and full sample strengthens the interpretation that the differences between treated and control groups reflect structural features of the data rather than sample-specific artifacts, and further justifies the inclusion of appropriate controls in the empirical specification.

Indeed, all specifications control for the log of GDP for both origin and partner countries, log of environmental taxes, and bilateral harmonic distance.

Aggregating the group-time ATTs into a single summary measure, the staggered estimator yields an overall ATT on log imports of 0.288 and on log exports of 0.068. In economic terms, the import ATT corresponds to bilateral imports between EU ETS members being on average about 33% higher relative to non-ETS partner pairs, while the export ATT is statistically noisy and economically zero. The import effect is therefore substantially larger than the baseline DiD estimate of approximately 20% (Table 3); this is consistent with the staggered-DiD design giving more weight to long-horizon differences and to the additional cohorts of later-entrant countries (Bulgaria and Romania in 2006, Iceland and Norway in 2007, Croatia in 2012). The export null is also reinforced: the point estimate is small and statistically insignificant, and the confidence interval comfortably includes zero. The export event-study plot displays noticeably wider confidence intervals than the import plot, particularly at long horizons, where only the 2000 cohort contributes to the average treatment effect and the not-yet-treated control set effectively collapses to the rest of the world.

5.4 Sectoral heterogeneity and the mechanism

The aggregate estimates leave open whether the premium reflects a mechanism specific to the regulated sectors or general bilateral integration between ETS countries, the single market and the 2004 enlargement included. The sectoral evidence is particularly informative because it allows us to distinguish between regulatory alignment and broader integration effects. A mechanism operating through common carbon pricing should be concentrated in ETS-covered sectors, while general integration between the same countries should affect covered and non-covered activities in a similar way.[16] We test this on a single supply-chain pair where treatment status is unambiguous and the horizontal combustion clause is least binding: basic iron and steel (ISIC 2710), covered since Phase I and accounting for roughly 12 per cent of verified ETS emissions (Mendelevitch et al., 2024), as the ETS sector, and fabricated metal products (ISIC 2811–2830), structurally below the 20 MW thermal-input threshold, as the non-ETS sector. This comparison is relevant because the two sectors share final buyers, trade geography, and market conditions. At the same time, they differ in one key dimension: carbon pricing applies directly to the production of basic iron and steel but not to fabricated metal products. Sector-level bilateral trade flows are constructed from BACI product-level data.[17]

[16] See Section 3.3.

[17] We use this specification because defining a clean non-ETS comparison group requires care. Annex I of Directive 2003/87/EC covers not only the explicitly listed activities, but also any installation with rated thermal input exceeding 20 MW regardless of sector (the horizontal combustion clause). To sharpen identification, we focus on a single supply-chain pair where treatment status is unambiguous and the combustion-clause concern is most easily addressed. A systematic, economy-wide analysis of sectoral heterogeneity, with a research design tailored specifically to that question, lies beyond the scope of this paper and is left for future research.

	imports (USD)		imports (tons)		exports (USD)		exports (tons)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ETS	non-ETS	ETS	non-ETS	ETS	non-ETS	ETS	non-ETS
DiD	0.345***	0.191*	0.428***	0.0879	0.417***	-0.223**	0.545***	-0.221*
	(0.113)	(0.108)	(0.146)	(0.120)	(0.154)	(0.110)	(0.178)	(0.126)
Obs	10365	10232	10360	10226	10477	10290	10476	10286
FE pair	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors clustered at the directed country-pair level in parentheses. Sample: top-10% bilateral partners. ETS sector: basic iron and steel (ISIC 2710), covered since 2005. Non-ETS sector: fabricated metal products (ISIC 2811–2830), downstream transformation, no own smelting furnaces. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Mechanism test: basic iron and steel vs fabricated metal products (top-10% partners).

Table 5 reports the results. Among ETS partners, bilateral imports of basic iron and steel, covered since Phase I, are about 41% higher (a coefficient of 0.345, significant at the 1% level), whereas for the otherwise comparable downstream sector of fabricated metal products the premium is markedly smaller and only marginally significant. The two sectors share buyers, trade geography and price signals, and differ only in whether the carbon price applies at the production stage; the contrast is therefore hard to square with general bilateral integration, which would lift covered and non-covered trade alike, and instead locates the effect in the activities the scheme actually regulates, as the alignment mechanism requires. The gap is sharper in quantity than in value terms.

On the export side the two sectors move in opposite directions: steel exports between ETS partners are roughly 50% higher, while fabricated-metal exports are some 20% lower. This pattern resolves the aggregate export null documented above: because downstream, non-covered goods account for the bulk of exports, the positive premium on regulated products and the negative effect on unregulated ones offset once summed, so that the aggregate conceals two countervailing forces that only the disaggregated data reveal.

The sectoral evidence rules out a pure integration story, under which trade among EU members would have responded alike in covered and non-covered sectors. We identify the mechanism from a single supply chain rather than a broad cross-section because cleanly delimiting non-ETS sectors is itself difficult: under the horizontal combustion clause of Annex I to the Directive, any installation above the thermal-input threshold falls within the scheme regardless of its nominal sector, blurring the covered/non-covered boundary elsewhere in the economy.

6 Robustness

We now show that the baseline import premium survives to alternative specifications and robustness checks. The single most important alternative, that the premium reflects general EU integration rather than carbon regulation, is addressed by the within-sector evidence in Section 5.4; here we take up the remaining concerns.

6.1 Common currency

A first concern is that the estimated premium reflects the introduction of the euro rather than EU ETS, given the substantial overlap between euro-area and EU ETS membership.[18] The evidence on the euro's own trade effect is itself mixed, ranging from sizeable (Rose, 2000; Rose and Van Wincoop, 2001) to negligible (Santos Silva and Tenreyro, 2010; Figueiredo et al., 2016). To separate the two channels, we re-estimate the baseline on the subsample of pairs in which at least one country has not adopted the euro.[19]

As shown in Figures 11, 12 and the Table 6, the intra-ETS import premium is slightly larger and remains significant once jointly euro-adopting pairs are excluded, so it does not reflect common-currency dynamics. Appendix D reports the complementary estimates

[18]Euro-area early adopters of the EU ETS include Austria, Belgium, Cyprus, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Slovakia and Spain. Estonia, Latvia, Lithuania and Slovenia adopted the euro later (2011, 2014, 2015 and 2007), and Malta in 2008.

[19]Non-euro EU ETS early adopters are the Czech Republic, Denmark, Hungary, Poland, Sweden and the United Kingdom.

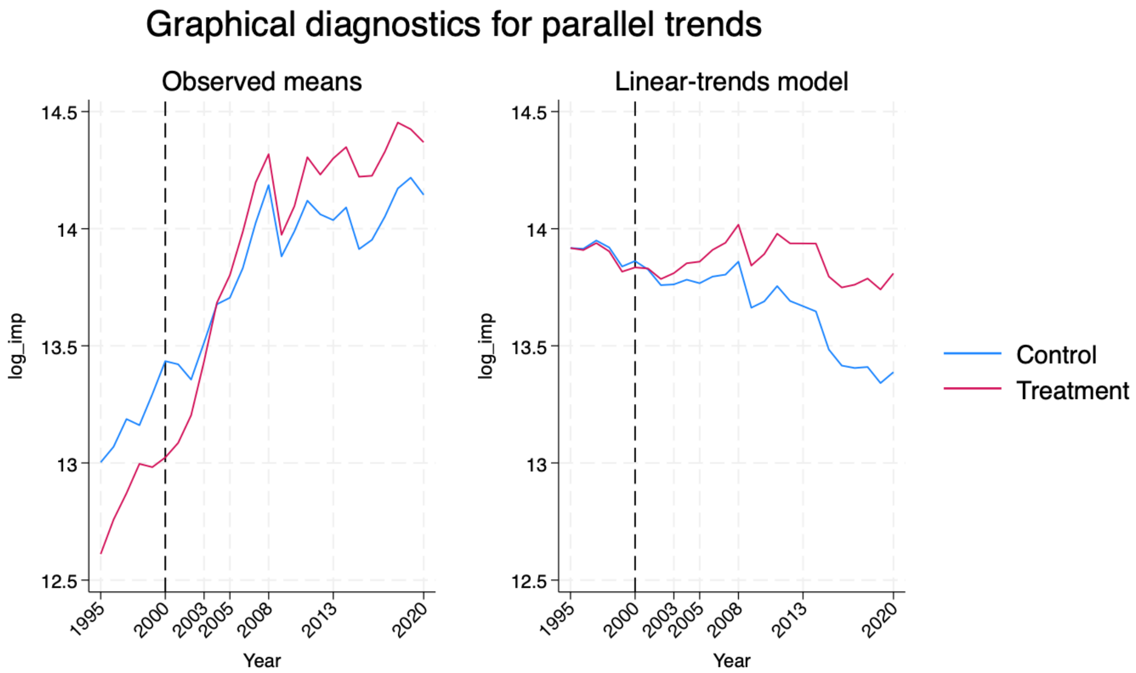

Figure 11: Log import baseline of non-euro country couples.

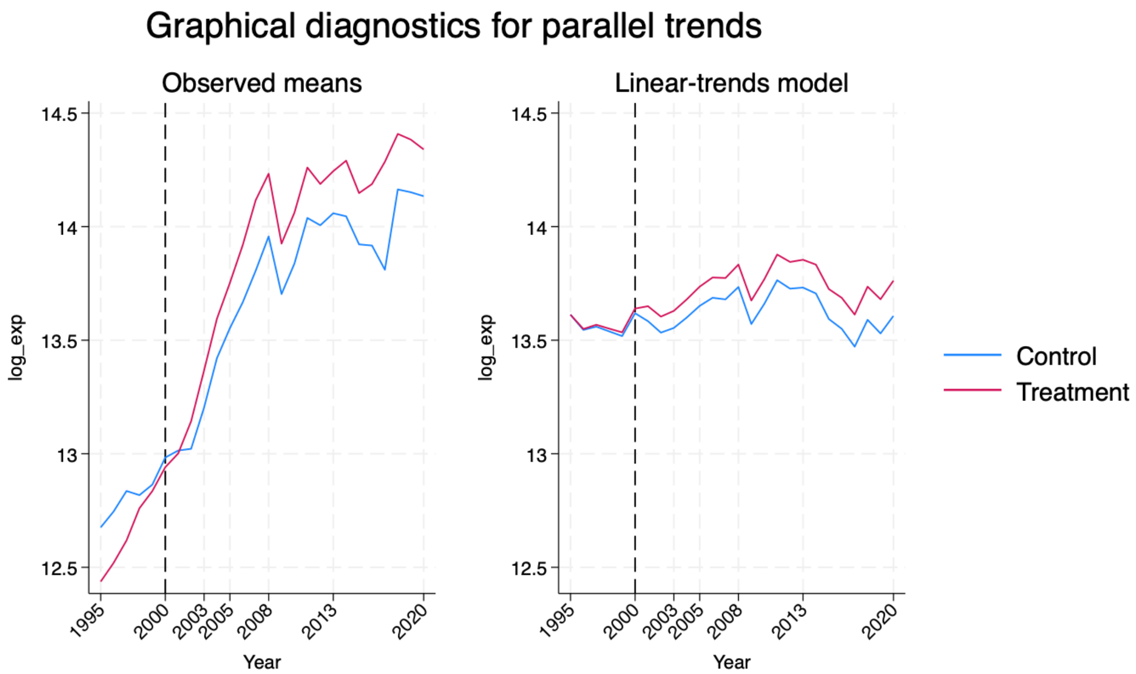

Figure 12: Log export baseline of non-euro country couples.

for pairs in which both countries adopt the euro.

	(1) log(imports)	(2) log(exports)
did	0.221***	0.0983
	(3.63)	(1.36)
$\log(\text{GDP}_{country})$	1.079***	0.570**
	(4.76)	(2.98)
$\log(\text{GDP}_{partner})$	0.851***	0.970***
	(4.01)	(5.04)
log(distance)	-14.93**	-6.905
	(-2.91)	(-1.25)
$\log(\text{env tax}_{partner})$	0.436***	0.189**
	(6.60)	(2.77)
$\log(\text{env tax}_{country})$	0.188***	0.450***
	(4.23)	(8.62)
FE year	✓	✓
FE id	✓	✓
Observations	8864	9113

t statistics in parentheses; standard errors clustered at the country-pair level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Baseline results of non-euro country couples.

6.2 Non-independence of mirrored trade flows

Because the origin dimension covers mirrored trade flows between the 30 EU ETS countries, every ETS–ETS pair enters the panel in both directions (for example, both Italy–Germany and Germany–Italy), and the two directions are nearly mirror images of one another. The baseline standard errors are clustered at the directed country-pair level, which treats the two directions as separate clusters. To account for the possibility that mirrored observations are not independent, we re-estimate the baseline specification clustering instead at the undirected-dyad level, so that the two directions of each pair share a common cluster. As shown in Table 7, the estimated coefficients are unchanged and their precision is virtually identical: the intra-ETS import premium remains significant at the 1% level, and the export effect remains small and insignificant. This stability reflects the fact that the directed country-pair fixed effects already absorb most of the within-dyad correlation, so that further aggregating the clusters has a negligible effect on inference.

	(1)	(2)	(3)	(4)
	log(imp)	log(exp)	log(imp)	log(exp)
	cl. pair	cl. dyad	cl. pair	cl. dyad
did	0.184***	0.184***	0.0482	0.0482
	(0.0561)	(0.0572)	(0.0660)	(0.0670)
Clusters	451	365	462	370
Observations	11050	11050	11365	11365

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

All specifications include gravity controls, country-pair and year fixed effects.

Table 7: Baseline estimates with alternative clustering: directed country-pair vs. undirected dyad.

6.3 Sensitivity to the top-trading-partner filter

A natural concern is that the top-10% filter inflates the estimate by over-representing large, well-integrated pairs. To address it, we re-estimate the baseline on the full set of trading partners, without the filter, and progressively relax the threshold (Table 8). Within log-OLS the import premium remains positive and significant at every threshold, while the export coefficient stays small and insignificant on the filtered samples.

	log(import)				log(export)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Top 10%	Top 25%	Top 50%	All	Top 10%	Top 25%	Top 50%	All
did	0.184***	0.126***	0.295***	0.484***	0.0482	-0.0316	0.0464	0.130***
	(0.0561)	(0.0445)	(0.0456)	(0.0468)	(0.0660)	(0.0479)	(0.0419)	(0.0401)
N	11050	24002	39926	51347	11365	23590	38431	50253

Standard errors in parentheses

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

Table 8: Sensitivity to the trading-partner threshold: imports and exports.

Table 8 shows that removing the filter does not shrink the import premium: the coefficient rises from 0.184 (about 20%) on the top-10% sample to 0.484 (about 62%) on the full set of partners. Far from inflating the effect, the filter therefore yields a more conservative estimate. This larger unfiltered coefficient should be read with caution. On the top-10% sample the parallel-trends test is satisfied ($F(1, 450) = 0.20$, $p = 0.65$), whereas on the full sample it is strongly rejected, so including marginal relationships introduces differential pre-trends that the filtered sample does not display. Section 6.4

reinforces this reading: under PPML, which corrects the heteroskedasticity bias of log-OLS, the import premium is about 11% whether or not the filter is applied, confirming that the unfiltered jump to 62% is an estimation artefact rather than a larger effect.

The export margin tells a consistent story. On the filtered sample the export effect is small and statistically insignificant, matching the aggregate null of Section 5. Once the filter is removed, the export coefficient turns positive and significant (0.130, about 14%). This apparent export premium, however, mirrors the import pattern in being confined to the unfiltered sample, where the parallel-trends test is again rejected. We therefore do not interpret it as a genuine reorientation of exports: like the inflated import coefficient, it reflects the differential pre-trends introduced by marginal partners rather than a treatment effect.

The filter delivers conservative magnitudes, and the qualitative results (a positive and significant import premium, no aggregate export effect) are those obtained on the sample where the identifying assumption is met.

6.4 Poisson Pseudo-Maximum Likelihood estimation (PPML)

Table 9 reports the baseline difference-in-differences estimated by Poisson pseudo-maximum likelihood, which estimates gravity in multiplicative form and is consistent under heteroskedasticity (Santos Silva and Tenreyro, 2006).[20] The import premium remains positive and significant: a coefficient of 0.109 implies intra-ETS imports about 11% higher than trade with non-participants, a more conservative magnitude than the log-OLS (20%) and staggered (33%) estimates. The export coefficient is small, negative and statistically insignificant, in line with the aggregate null.

Table 10 re-estimates the same specification on all trading partners, without the top-10% filter. The import premium is essentially unchanged (0.104, again about 11%)

[20] An alternative route to the zero and small flows, and to the selection of pairs into positive trade, is the two-stage selection estimator of Helpman et al. (2008), which models selection into trading partners explicitly. In our top-10% sample, where zeros are essentially absent, the two approaches should largely coincide on the intensive margin.

and the export coefficient remains insignificant. This stability is informative: the large unfiltered log-OLS coefficient (about 62%, Table 8) does not survive an estimator that accommodates the small flows and heteroskedasticity of log-linearized gravity, confirming that it reflects estimation bias rather than a genuinely larger effect. Across estimators and samples the qualitative conclusion is the same: a positive, significant import premium of about 11 to 20% and no effect on exports.

	(1)	(2)
	Imports (PPML)	Exports (PPML)
did	0.109**	-0.0307
	(0.0478)	(0.0513)
Observations	12,216	12,562

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: PPML results for imports and exports.

	(1)	(2)
	Imports (PPML, all)	Exports (PPML, all)
did	0.104**	-0.0124
	(0.0442)	(0.0449)
Observations	54,552	53,445

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: PPML results for imports and exports without filter.

7 Conclusions

This paper asks whether common carbon pricing reshapes the selection of trading partners. Embedding a gravity model in a difference-in-differences design, we examine whether countries subject to the same carbon-pricing regime trade differently from pairs in which one partner lies outside it.

Our main result is that EU ETS participation raises bilateral imports among participating countries relative to trade with non-participants, by about 20% in the baseline and about 33% in the staggered design of Callaway and Sant'Anna (2021), and a more conservative, sample-stable estimate of about 11% under PPML, while the export effect is not statistically significant. The import premium is stable across the baseline, event-study and staggered specifications, survives the removal of the top-10% partner filter (where the larger log-OLS coefficient is an estimation artefact that PPML removes), is unaffected by clustering on the undirected dyad, and persists when jointly euro-adopting pairs are excluded, so it does not merely reflect monetary integration. The alternative most likely explanation, that the estimated premium reflects broader European integration rather than carbon regulation, is difficult to reconcile with the sector-level evidence: the import premium is markedly stronger in activities directly covered by the EU ETS than in otherwise comparable non-covered sectors, where on the export margin the two even move in opposite directions, a pattern a uniform integration effect would not produce.

We read these findings through the two channels of Section 3. Carbon pricing raises the marginal cost of regulated firms, which on its own would divert trade away from regulated economies, and it reduces regulatory asymmetries between co-regulated partners, which on its own encourages trade within the regulated group. Because the two forces have opposite signs, our reduced-form coefficient does not test a signed prediction; it measures their net balance, and the positive import estimate says that, on the sourcing margin, the alignment channel dominates. The aggregate export null does not reflect an absence of effects: the sector-level analysis shows that it masks a positive premium in the regulated sector and a negative effect downstream that offset once aggregated. The evidence also satisfies the three restrictions the framework places on the data regardless of which channel dominates: the effect is absent before the policy and emerges only afterwards, it strengthens as allowance prices rise across phases, and it concentrates in the sectors the EU ETS actually regulates while being absent in an otherwise comparable non-covered sector (Section 5.4). The last is crucial, since general integration would have moved covered and non-covered sectors alike.

The results are consistent with the EU ETS operating as a de facto climate club

(Nordhaus, 2015): a common carbon-pricing regime strengthens trade integration among insiders, on the sourcing margin, even without explicit penalties on outsiders, plausibly because firms favour suppliers under comparable carbon constraints and anticipate future border-adjustment measures such as the CBAM (Handley and Limão, 2015, 2017; Chen, 2023). Our results are not a comprehensive test of carbon leakage: a relative shift toward co-regulated partners does not rule out leakage at finer levels of disaggregation or in embodied-emission terms. What they show is that, at the aggregate bilateral level, trade does not reorient toward unregulated partners as the dominant response to carbon pricing.

The implications for policy are twofold. First, the trade effects of climate policy cannot be read through cost increases alone: regulatory alignment is a quantitatively relevant channel that can offset, and on the import margin reverse, the leakage incentives created by asymmetric carbon pricing. Second, the climate-club logic suggests that broader coordination of carbon-pricing institutions may deliver trade-integration benefits alongside the standard environmental rationale for cooperation. This trade-integration benefit operates in the opposite direction to the transition costs and weak participation incentives emphasised by cost-focused analyses such as Ernst et al. (2023); our reduced-form evidence documents the benefit but is silent on the net welfare balance. The sector-level evidence reinforces this reading, since the premium is concentrated where the EU ETS actually bites. Natural extensions include assessing how the CBAM and the linkage of the EU ETS with other systems amplify the alignment premium, and extending the analysis to embodied-emission flows so as to gauge leakage in physical rather than value terms.

A Appendix: gravity model

The gravity model of international trade, initially formulated by Tinbergen (1962), represents one of the most empirically validated frameworks in the analysis of bilateral trade flows. Analogous to Newton's law of universal gravitation, the model posits that trade volume between two countries is proportional to the product of their economic sizes and inversely proportional to the geographic distance between them:

$$T_{ij} = A \cdot \frac{Y_i Y_j}{D_{ij}} \tag{4}$$

where T_{ij} denotes trade from country i to country j, Y_i and Y_j are the GDPs of the respective countries, D_{ij} is the distance between them, and A is a constant capturing other trade-affecting factors. The log-linearized form, suitable for empirical estimation, is:

$$\ln(T_{ij}) = \ln(A) + \alpha \ln(Y_i) + \beta \ln(Y_j) - \theta \ln(D_{ij}) + \varepsilon_{ij} \tag{5}$$

Here, θ represents the elasticity of trade with respect to distance and typically assumes a negative value, indicating a trade-reducing effect as distance increases. While globalization has reduced transport and information costs, empirical findings consistently report a relatively stable and significantly negative distance coefficient, a robust pattern often referred to as the "distance puzzle".

The theoretical foundation of the gravity model was later strengthened by Anderson (1979), and formalized by Anderson and Van Wincoop (2003), who introduced multilateral resistance terms. The extended model includes country-specific factors and bilateral trade frictions:

$$E[x_{ij}] = A_i^X A_j^M D_{ij}^{-\theta} \exp(\lambda L_{ij}) \tag{6}$$

where L_{ij} captures additional dyadic factors such as common language or colonial ties, and λ is the associated coefficient. The log-linear form becomes:

$$\ln(x_{ij}) = \ln(A_i^X) + \ln(A_j^M) - \theta \ln(D_{ij}) + \lambda L_{ij} + \ln(\eta_{ij}) \tag{7}$$

Despite its robustness, the magnitude of θ varies across studies. To explore this variation, Disdier and Head (2008) conducted a meta-analysis aggregating 1,467 estimates of the distance coefficient from 103 empirical studies. Their findings show an average value of $\theta \approx 0.91$, implying that a 10% increase in distance reduces trade by about 9.1%. This is consistent with earlier empirical work, where estimates typically fall between 0.6 and 1.5 in absolute value.

The authors develop a regression to explain the diversity of distance effects across different studies. The model used is:

$$\hat{\theta}_{ij} = u_i + \beta X_{ij} + e_{ij}$$

Where $\hat{\theta}_{ij}$ is the estimated distance coefficient for the j-th observation in study i, u_i represents the random effects associated with the studies, X_{ij} is the matrix of independent variables explaining the variability of distance effects, and e_{ij} is the residual error. The authors find that the effect of distance on trade is negative and has been increasing since the 1970s.

For a more precise analysis, they define four temporal specifications: before 1970, the 1970s, the 1980s, and the period from the 1990s onwards. They find that in the most recent period, trade has decreased by 37% compared to the period before 1970. The authors also perform controls for various instrumental variables, such as “single continent” and “developed countries,” finding that intra-continental trade is more elastic than intercontinental trade (distance has a greater impact on land transport costs) and that in more developed countries, the distance effect is weaker due to better infrastructure. “Common language” significantly increases the effect of distance, while “membership in a trade agreement” also has a positive impact, though smaller than the other controls.

B Appendix: carbon price in the EU ETS

Figure 13: Historical price evolution of EU allowances. Source: Homaio.

C Appendix: 2005 Results

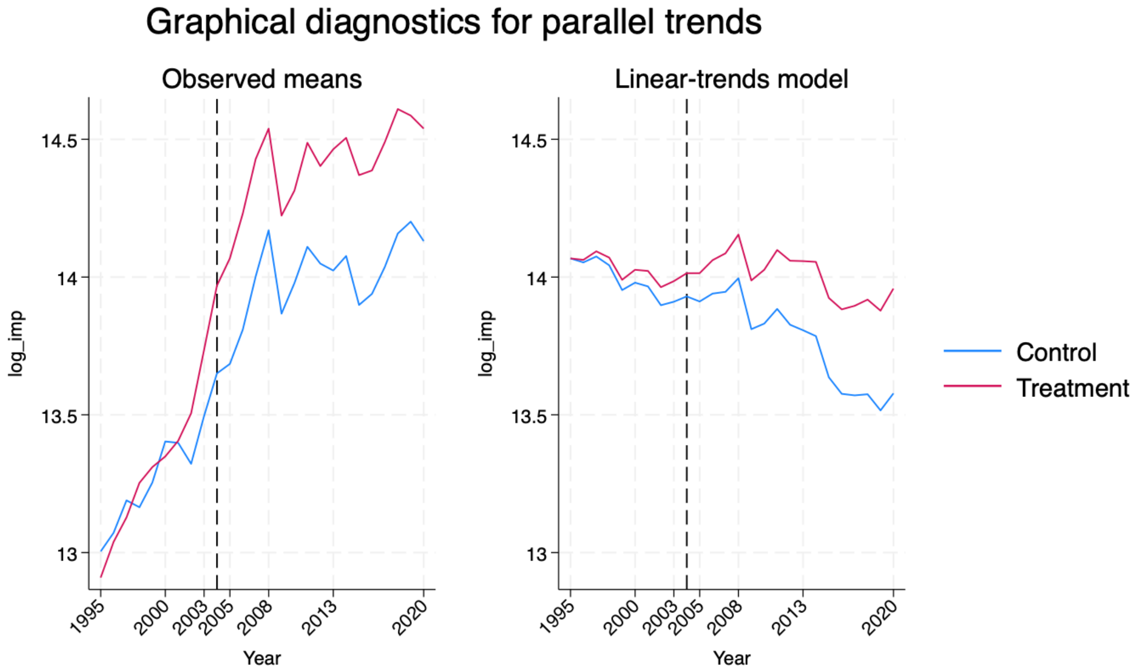

Figure 14: Baseline sample: log import in treated and control group.

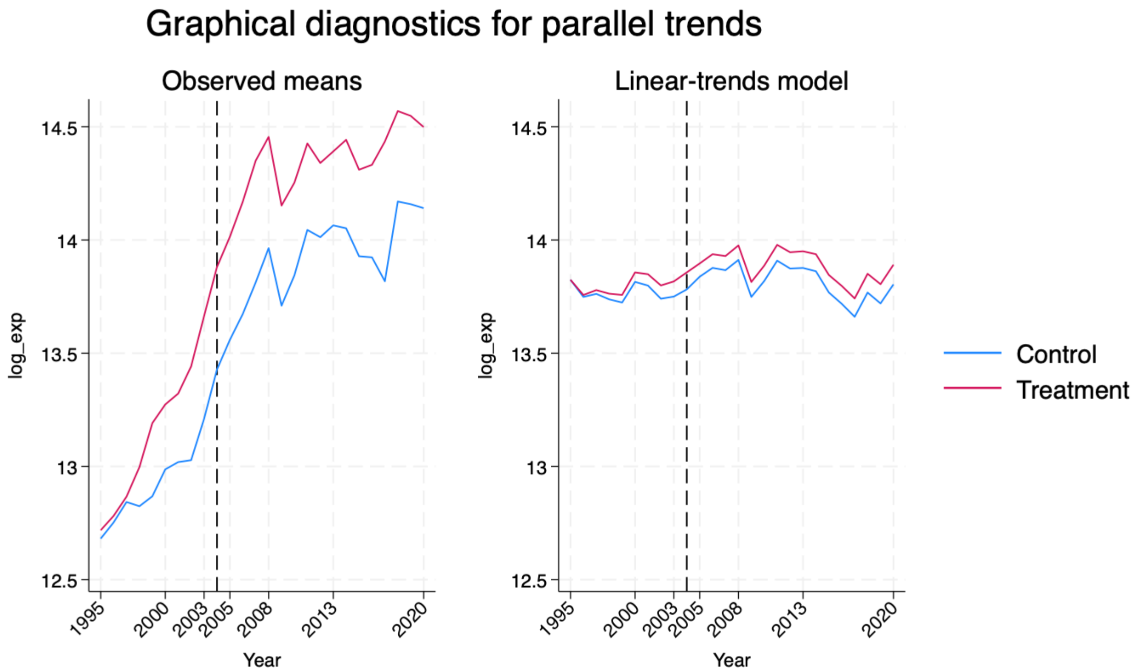

Figure 15: Baseline sample: log export in treated and control group.

D Appendix: euro vs non-euro

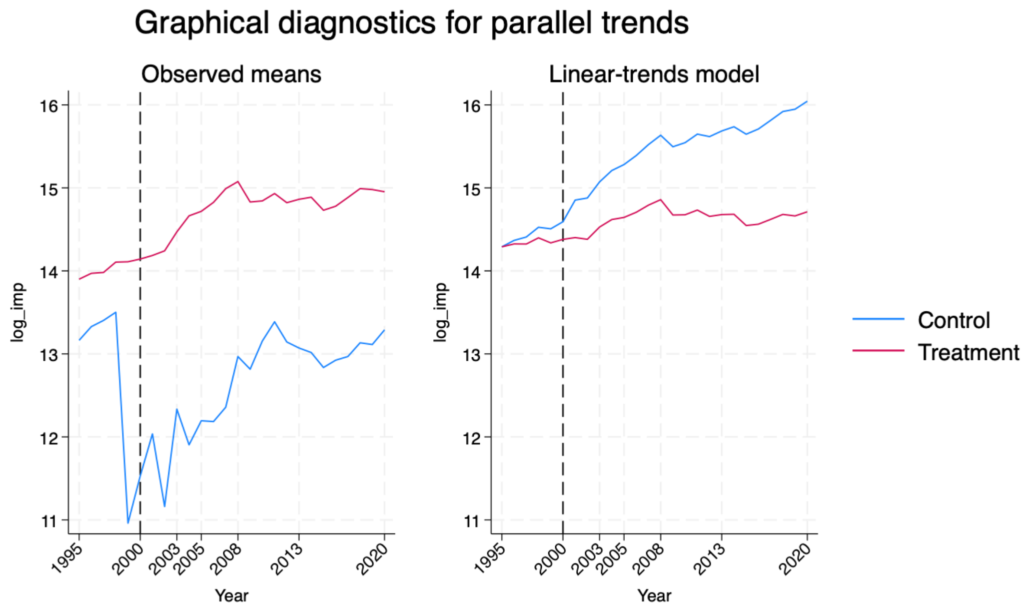

Figure 16: log import baseline of euro country couples - 2000

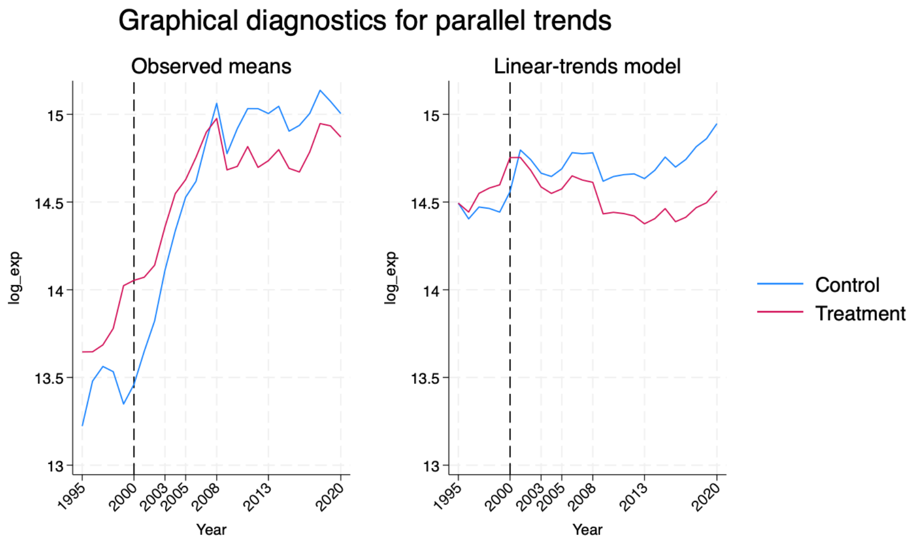

Figure 17: log export baseline of euro country couples - 2000

E Appendix: EU ETS emission share

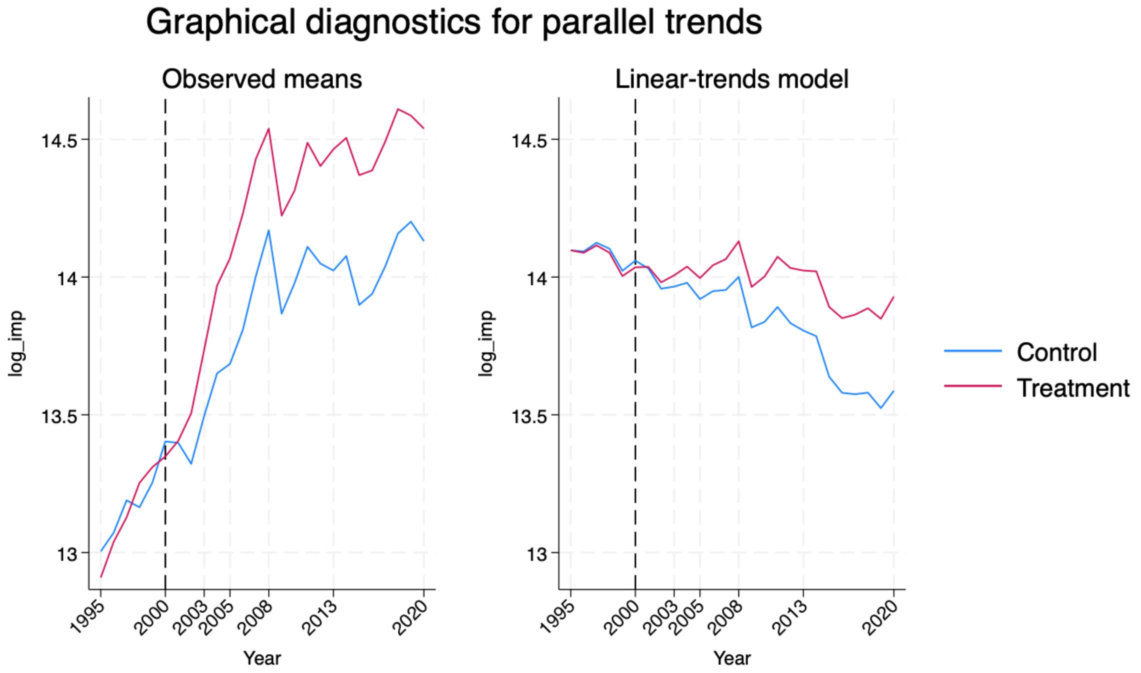

Figure 18: Baseline log import controlling for the share of emissions under the EU ETS.

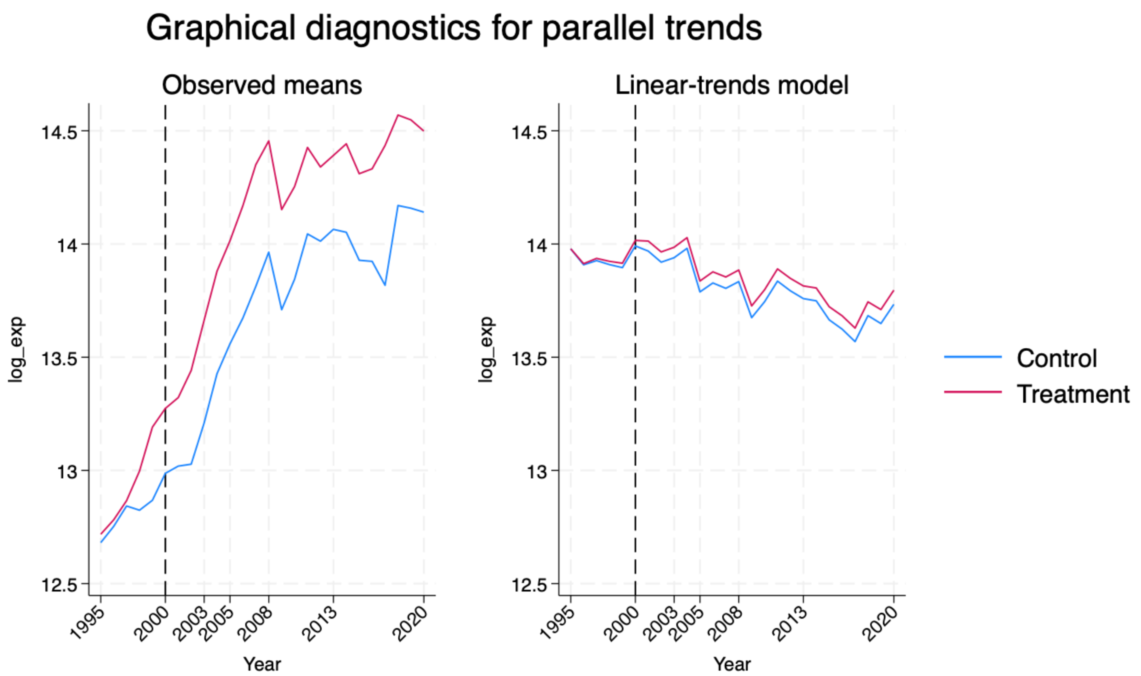

Figure 19: Baseline log export with for the share of emissions under the EU ETS.

References

Aichele, R. and G. Felbermayr (2015). Kyoto and carbon leakage: An empirical analysis of the carbon content of bilateral trade. *Review of Economics and Statistics 97*(1), 104–115.

Ambec, S., M. A. Cohen, S. Elgie, and P. Lanoie (2013, January). The Porter hypothesis at 20: Can environmental regulation enhance innovation and competitiveness? *Review of Environmental Economics and Policy 7*(1), 2–22.

Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *The American Economic Review 69*(1), 106–116.

Anderson, J. E. and E. Van Wincoop (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review 93*(1), 170–192.

Antoci, A., S. Borghesi, G. Iannucci, and M. Sodini (2021). Should I stay or should I go? carbon leakage and ets in an evolutionary model. *Energy Economics 103*, 105561.

Baier, S. L. and J. H. Bergstrand (2007). Do free trade agreements actually increase members' international trade? *Journal of International Economics 71*(1), 72–95.

Bernard, J.-T. and M. Kichian (2021). The impact of a revenue-neutral carbon tax on GDP dynamics: The case of British Columbia. *The Energy Journal 42*, 27–33.

Biancalani, F., G. Gnecco, R. Metulini, and M. Riccaboni (2024). The impact of the European Union emissions trading system on carbon dioxide emissions: a matrix completion analysis. *Scientific Reports 14*, 19676.

Branger, F., P. Quirion, and J. Chevallier (2016). Carbon leakage and competitiveness of cement and steel industries under the EU ETS: much ado about nothing. *The Energy Journal 37*(3), 109–136.

Bremer, L. and K. Sommer (2025). Economic performance and investments under emissions trading: Untangling the effects of a staggered regulation. *Energy Economics 142*, 108170.

Calel, R. and A. Dechezleprêtre (2016). Environmental policy and directed technological change: evidence from the European carbon market. *Review of economics and statistics 98*(1), 173–191.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Chen, G. (2023). Impact of carbon border adjustment mechanism on china's manufacturing sector: a dynamic recursive cge model based on an evolutionary game. *Journal of Environmental Management 347*, 119029.

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

De Beule, F., F. Schoubben, and K. Struyfs (2022). The pollution haven effect and investment leakage: The case of the EU-ETS. *Economics Letters 215*, 110536.

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

Dechezleprêtre, A. and M. Sato (2017). The impacts of environmental regulations on competitiveness. *Review of Environmental Economics and Policy 11*(2), 183–206.

Disdier, A.-C. and K. Head (2008). The puzzling persistence of the distance effect on bilateral trade. *The Review of Economics and Statistics 90*(1), 37–48.

Ellerman, A. D. and B. K. Buchner (2008). Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005–06 emissions data. *Environmental and Resource Economics 41*(2), 267–287.

Ernst, A., N. Hinterlang, A. Mahle, and N. Stähler (2023). Carbon pricing, border adjustment and climate clubs: Options for international cooperation. *Journal of International Economics 144*, 103772.

EU Parliament (2003). Directive 2003/87/ec. *Official Journal of European Union*.

European Commission (2009). EU Action Against Climate Change: The EU Emissions Trading Scheme (EU ETS).

European Commission (2020). EU ETS Handbook (Updated 2020).

Falk, M. and E. Hagsten (2019). Short-run impact of the flight departure tax on air travel. *International Journal of Tourism Research 21*(1), 37–44.

Figueiredo, E., L. R. Lima, and G. Schaur (2016). The effect of the euro on the bilateral trade distribution. *Empirical Economics 50*(1), 17–29.

Fowlie, M. and M. Reguant (2018). Challenges in the measurement of leakage risk. *AEA Papers and Proceedings 108*, 124–29.

Handley, K. (2014). Exporting under trade policy uncertainty: Theory and evidence. *Journal of International Economics 94*(1), 50–66.

Handley, K. and N. Limão (2015). Trade and investment under policy uncertainty: Theory and firm evidence. *American Economic Journal: Economic Policy 7*(4), 189–222.

Handley, K. and N. Limão (2017). Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States. *American Economic Review 107*(9), 2731–2783.

Helpman, E., M. Melitz, and Y. Rubinstein (2008). Estimating trade flows: Trading partners and trading volumes. *The Quarterly Journal of Economics 123*(2), 441–487.

Hintermann, B. and M. Ludwig (2023). Home country bias in international emissions trading: Evidence from the EU ETS. *Resource and Energy Economics 71*, 101336.

Känzig, D. R. (2023). The unequal economic consequences of carbon pricing. NBER Working Papers 31221, National Bureau of Economic Research, Inc.

Konradt, M. and B. Weder (2021). Carbon taxation and inflation: Evidence from the european and canadian experience. Technical Report HEIDWP17-2021, Graduate Institute of International and Development Studies Working Paper.

Larch, M. and J. Wanner (2017). Carbon tariffs: An analysis of the trade, welfare, and emission effects. *Journal of International Economics 109*, 195–213.

Mangiante, G. (2024). The geographic effects of carbon pricing. *European Economic Review 168*, 104825.

Mendelevitch, R., H. Hermann, V. Graichen, T. Bibu, F. Lettow, and C. Nissen (2024). Development of the iron and steelmaking sector under the EU ETS. *German Environment Agency*.

Metcalf, G. E. (2019). On the economics of a carbon tax for the united states. *Brookings Papers on Economic Activity*.

Metcalf, G. E. and J. H. Stock (2023). The macroeconomic impact of Europe's carbon taxes. *American Economic Journal: Macroeconomics 15*(3), 265–86.

Naegele, H. and A. Zaklan (2019). Does the EU ETS cause carbon leakage in European manufacturing? *Journal of Environmental Economics and Management 93*, 125–147.

Narassimhan, E., K. S. Gallagher, S. Koester, and J. R. Alejo (2018). Carbon pricing in practice: A review of existing emissions trading systems. *Climate Policy 18*(8), 967–991.

Nordhaus, W. (2015). Climate clubs: Overcoming free-riding in international climate policy. *American Economic Review 105*(4), 1339–70.

Oesingmann, K. (2022). The effect of the European emissions trading system (EU ETS) on aviation demand: An empirical comparison with the impact of ticket taxes. *Energy Policy 160*, 112657.

Pacelli, A. (2025). Being a new entrant or an incumbent matters: emission abatement and the eu ets. Technical report.

Porter, M. E. and C. v. d. Linde (1995). Toward a new conception of the environment-competitiveness relationship. *Journal of economic perspectives 9*(4), 97–118.

Rose, A. K. (2000). One money, one market: the effect of common currencies on trade. *Economic policy 15*(30), 8–45.

Rose, A. K. and E. Van Wincoop (2001). National money as a barrier to international trade: The real case for currency union. *American economic review 91*(2), 386–390.

Santos Silva, J. and S. Tenreyro (2010). Currency unions in prospect and retrospect. *Annu. Rev. Econ. 2*(1), 51–74.

Santos Silva, J. M. C. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics 88*(4), 641–658.

Tinbergen, J. (1962). *Shaping the world economy; suggestions for an international economic policy*. New York : Twentieth Century Fund.

Wang, M. and T. Kuusi (2024). Trade flows, carbon leakage, and the EU emissions trading system. *Energy Economics 134*, 107556.

Woerdman, E., A. Arcuri, and S. Clò (2008). Emissions trading and the polluter-pays principle: Do polluters pay under grandfathering? *Review of Law and Economics 4*(2), 565–590.