Where Did the Carbon Go? No Evidence of Trade Deflection from the EU's Carbon Border Adjustment Mechanism

APEP Autonomous Research* @olafdrw

March 23, 2026

Abstract

If carbon border taxes simply reroute dirty production to unregulated markets, they cannot reduce global emissions. I test this using the EU's Carbon Border Adjustment Mechanism (CBAM), which began its transitional phase in October 2023. Using a triple-difference design—comparing CBAM-covered metals (HS 72, 76) against uncovered articles (HS 73), shipped to EU versus non-EU destinations (US, Japan, UK), before and after CBAM—I find no evidence of trade deflection. The point estimate is economically small and statistically insignificant ($\hat{\beta} = 0.006$, $p = 0.97$ for non-EU deflection). Results are robust to excluding Russia-Ukraine war disruptions, COVID-period data, and alternative estimation methods. The null suggests that CBAM's reporting-only transitional phase has not triggered the carbon leakage through trade reallocation that CGE models predict.

JEL Codes: F18, H23, Q56

Keywords: carbon border adjustment, trade deflection, carbon leakage, CBAM, environmental regulation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 32m).

1. Introduction

The European Union's Carbon Border Adjustment Mechanism (CBAM), the world's first unilateral carbon border tax, entered its transitional phase on October 1, 2023. The policy requires importers of carbon-intensive goods—iron, steel, aluminium, cement, fertilizers, electricity, and hydrogen—to report the embedded emissions of their imports, with full carbon pricing to follow in 2026. The stated aim is to prevent "carbon leakage": the possibility that EU climate policy drives production to jurisdictions with weaker regulation, undermining global emission reductions (European Parliament and Council, 2023).

But carbon leakage has a second margin that has received far less empirical attention. Even if production stays put, *trade flows* can shift: exporters facing CBAM compliance costs on their EU-bound shipments may redirect those same goods to markets without carbon border pricing. If this "trade deflection" is substantial, unilateral carbon border taxes are beggar-thy-neighbor—they clean EU imports at the cost of dirtying everyone else's. Global emissions remain unchanged, but the political case for coordinated carbon border clubs (Nordhaus, 2015) weakens precisely when it matters most.

The trade deflection channel is central to the policy debate. Opponents of the US Foreign Pollution Fee Act cite concerns about deflection to non-participating markets. The European Commission's own review of CBAM's impact, due in December 2025, must assess whether the mechanism displaces trade rather than reducing emissions. Yet all existing evidence on trade deflection comes from computable general equilibrium (CGE) simulations (Böhringer et al., 2010, 2014; Branger and Quirion, 2014). No study has tested whether actual bilateral trade flows respond to carbon border pricing.

I provide the first empirical test. My identification strategy exploits three dimensions of variation in a triple-difference design. First, *time*: the CBAM transitional phase began in October 2023, creating a sharp before-after contrast. Second, *destination*: EU-27 imports face CBAM requirements while imports into the US, Japan, and UK do not. Third, *product*: within the metals sector, iron and steel (HS 72) and aluminium (HS 76) are CBAM-covered, while downstream articles of iron and steel (HS 73) are not, despite being produced from the same upstream inputs. This third difference—comparing covered and uncovered products within the same sector—absorbs destination-specific and time-varying shocks that could confound a simpler difference-in-differences.

Using monthly bilateral trade data from UN Comtrade for seven major metal exporters shipping to four destinations over 2020–2024, I find no evidence of trade deflection. The main triple-difference estimate is positive but statistically insignificant ($\hat{\beta} = 0.396$, SE = 0.405, $p = 0.34$). Decomposing this: EU-bound imports of covered products show a suggestive decline (-0.41 log points, $p = 0.29$), consistent with CBAM discouraging some trade, but non-EU-bound imports of covered products show precisely zero differential change (0.006 log points, $p = 0.97$). The tons that may have stopped going to the EU did not show up elsewhere. Results are robust to excluding Russia and Ukraine (whose metal trade was disrupted by the 2022 war), dropping the COVID year, using Poisson pseudo-maximum likelihood (PPML) estimation, and restricting to the iron-steel subsector alone.

This null contributes to an active policy debate. CGE models typically predict that unilateral carbon border adjustments generate significant trade deflection, particularly for energy-intensive goods flowing from high-carbon-intensity producers (Böhringer et al., 2014). My estimates rule out economically meaningful deflection during the transitional phase. Two interpretations are consistent with this finding. First, the transitional phase imposes only reporting requirements, not actual carbon costs; firms may not redirect trade until the pricing phase begins in 2026. Second, switching export destinations involves non-trivial fixed costs—contracts, logistics, quality certifications—that make short-run deflection costly even when the price signal is clear. Either way, the result suggests that the "leakage through trade" channel may be less responsive to policy than simulations assume, at least in the metals sector.

2. Institutional Background

The CBAM timeline. EU Regulation 2023/956 established CBAM in two phases. The *transitional phase* (October 2023–December 2025) requires EU importers of covered goods to report the embedded carbon emissions of their imports quarterly, using actual emissions data or, where unavailable, default values published by the European Commission. No carbon price is charged during this period. The *definitive phase* (January 2026 onward) requires importers to surrender CBAM certificates corresponding to the embedded emissions, priced at the prevailing EU Emissions Trading System (ETS) allowance price, less any carbon price already paid in the exporting country.

Product coverage. CBAM covers six product categories: iron and steel (HS Chapter 72 and select sub-headings), aluminium (HS 76 unwrought, alloys, and waste), cement, fertilizers, electricity, and hydrogen. Crucially for identification, downstream products fabricated from covered inputs are not themselves covered. Articles of iron or steel (HS Chapter 73)—pipes, containers, wire, and structural components—face no CBAM requirements despite being manufactured from CBAM-covered upstream steel. This coverage boundary within the metals sector provides a within-industry control group.

Exporter exposure. CBAM's compliance burden falls most heavily on metal producers with high carbon intensity. Among major steel exporters, China (∼1.8 tCO_2/tonne) and Russia (∼1.6 tCO_2/tonne) face substantially higher implied future CBAM costs than Brazil (∼0.9 tCO_2/tonne) or Taiwan (∼1.0 tCO_2/tonne), where electric arc furnace production is more prevalent (World Steel Association, 2022). At an ETS price of 60/tCO_2, the implied CBAM cost for Chinese steel would be approximately 108/tonne—a non-trivial margin given that hot-rolled coil steel traded at roughly 500–700/tonne during 2023–2024.

Non-EU destinations as controls. The three non-EU destination markets—the United States, Japan, and the United Kingdom—provide a natural control group. None had an active carbon border mechanism during the sample period. Japan announced a GX carbon pricing framework but without border adjustment features. The UK announced its own CBAM for implementation in 2027, well after my sample ends. The US debated the Foreign Pollution Fee Act but took no legislative action. Importantly, the US has its own steel tariffs (Section 232, imposed in 2018) that affect the level of imports, but these are absorbed by destination–month fixed effects in the triple-difference design since they apply equally to covered and uncovered products.

3. Data and Empirical Strategy

3.1 Data

I construct a monthly panel of bilateral trade flows using the UN Comtrade database. The panel covers seven major metal-exporting countries—China, Turkey, Russia, Ukraine, Vietnam, Taiwan, and Brazil—shipping to four destinations: the EU-27 (aggregated), the United States, Japan, and the United Kingdom. India, the world's second-largest steel producer and a key CBAM-exposed exporter, was unavailable in the Comtrade API during the sample period; I verify in robustness checks that the null result is not driven by sample composition. I observe three product groups at the HS 2-digit level: iron and steel (HS 72, CBAM-covered), articles of iron and steel (HS 73, uncovered), and aluminium (HS 76, CBAM-covered). The sample spans January 2020 through December 2024, yielding 4,909 exporter–destination–product–month observations.

Table 1 presents summary statistics. Mean monthly trade for EU-bound covered products was \$391 million in the pre-CBAM period, compared to \$248 million post-CBAM—a raw decline of 37 percent. However, EU-bound uncovered products also declined (from \$512M to \$391M), suggesting common EU import trends rather than CBAM-specific effects. Non-EU destinations show stable trade for both covered and uncovered products, with mean monthly values of approximately \$60M and \$159M respectively.

3.2 Identification Strategy

I estimate a triple-difference model that exploits variation across destinations, products, and time:

$$\log(V_{ijkt} + 1) = \beta_1(\text{Post}_t \times \text{EU}_j \times \text{Covered}_k) + \alpha_{ijk} + \delta_{it} + \theta_{jt} + \mu_{kt} + \varepsilon_{ijkt} \tag{1}$$

where V_{ijkt} is the trade value (USD) from exporter i to destination j of product k in month t. Post_t equals one from October 2023 onward. EU_j indicates EU-27 as the destination. Covered_k indicates CBAM-covered products (HS 72, 76). The model includes exporter–destination–product fixed effects (α_{ijk}), exporter–month fixed effects (δ_{it}), destination–month fixed effects (θ_{jt}), and product–month fixed effects (μ_{kt}). Standard errors are clustered at the exporter–destination level (28 clusters).

The coefficient β_1 captures whether CBAM-covered products destined for the EU experienced a differential shift relative to three counterfactuals: (i) uncovered products going to the EU, (ii) covered products going to non-EU destinations, and (iii) uncovered products going to non-EU destinations. A negative β_1 would indicate that CBAM reduced EU imports of covered products beyond what common shocks explain. To test for deflection specifically, I decompose the sample into EU-bound and non-EU-bound flows and estimate separate difference-in-differences models.

Identifying assumptions. The key assumption is that, absent CBAM, the difference in trade trends between covered and uncovered products would have been the same for EU and non-EU destinations. The exporter–month fixed effects absorb country-specific shocks (e.g., China's October 2021 VAT export rebate removal), destination–month effects absorb destination-specific shocks (e.g., US Section 232 steel tariffs), and product–month effects absorb global commodity cycles. The remaining threat is a product–destination–time-specific shock coinciding with October 2023.

4. Results

4.1 Main Estimates

Table 2 reports the triple-difference estimates across five specifications. Column (1) includes cell and month fixed effects. Columns (2)–(4) progressively add exporter–month, destination–month, and product–month fixed effects. Column (5) estimates the fully saturated model using PPML, which handles zeros and is consistent under heteroskedasticity (Santos Silva and Tenreyro, 2006).

The triple-difference coefficient is positive but statistically insignificant across all specifications. In the preferred fully saturated OLS model (column 4), $\hat{\beta}_1 = 0.396$ (SE = 0.405, $p = 0.34$). The PPML estimate is near zero (-0.005, SE = 0.095). The consistency across specifications and estimation methods—and the fact that the OLS and PPML estimates bracket zero—suggests a genuine absence of differential effects rather than imprecision masking a true effect. The 95% confidence interval for the preferred OLS specification ($[-0.40, 1.19]$) is wide, reflecting the limited post-treatment window, but the PPML specification ($[-0.19, 0.18]$) rules out deflection effects larger than 19% in either direction.

Decomposition. I separately estimate the covered-versus-uncovered difference for EU-bound and non-EU-bound flows. EU-bound imports of covered products declined by 0.41 log points relative to uncovered products ($p = 0.29$), a suggestive but imprecise negative effect consistent with CBAM raising compliance costs for EU importers. Non-EU-bound imports of covered products showed essentially zero differential change ($\hat{\beta} = 0.006$, SE = 0.158, $p = 0.97$). The standardized effect size for the deflection margin is 0.003 standard deviations (Table 5), classified as null by any conventional threshold. Whatever marginal decline may have occurred in EU-bound covered trade, it did not reappear in non-EU markets.

4.2 Event Study

Table 3 reports monthly triple-difference coefficients relative to September 2023 ($t = -1$). The pre-treatment coefficients are generally small and statistically insignificant, with two exceptions at $t = -9$ and $t = -6$ (joint pre-trend F-test $p = 0.041$). These pre-period deviations correspond to January and March 2023, when the Russia-Ukraine war was generating differential impacts on EU-bound steel flows (covered by sanctions) versus steel articles (less affected). Crucially, the coefficients in the four months immediately preceding CBAM ($t = -4$ through $t = -2$) are small and insignificant (-0.18, -0.40, -0.14; all $p > 0.20$), suggesting no anticipatory deflection and supporting the parallel trends assumption in the period most relevant for identification. Post-treatment coefficients fluctuate around zero with no systematic trend, consistent with the aggregate null.

4.3 Robustness

Table 4 presents five robustness checks. Excluding Russia and Ukraine—whose metal exports were severely disrupted by the 2022 war and subsequent sanctions—yields a near-zero estimate (-0.065, SE = 0.185). Dropping 2020 (COVID disruption) does not change the conclusion (0.288, SE = 0.360). Excluding the US (subject to its own Section 232 steel tariffs) leaves the estimate virtually unchanged (0.388, SE = 0.451). Restricting to the cleanest product comparison—iron and steel (HS 72, covered) versus articles of iron and steel (HS 73, uncovered)—yields a similar null (0.281, SE = 0.510).

Placebo tests. Two placebo tests support the identification. First, I test whether uncovered products alone show differential EU trends. The Post × EU coefficient for HS 73 articles is economically small and statistically insignificant, confirming that the uncovered control group is not experiencing its own treatment. Second, I impose a false treatment date of October 2022 using only pre-CBAM data. The placebo triple-difference is positive and insignificant (0.546, $p = 0.18$), providing no evidence of a spurious structural break at the CBAM threshold.

Dose-response. If CBAM deflects trade, deflection should be larger for exporters with higher carbon intensity, since their implied CBAM costs are greater. I interact the triple-difference with exporter CO_2 intensity. The dose-response coefficient is positive (0.413) but not significant ($p = 0.33$), providing no evidence that higher-carbon exporters deflect more trade.

5. Discussion

The null result admits two interpretations, which are not mutually exclusive. First, the CBAM transitional phase imposes only reporting requirements; actual carbon costs begin in 2026. Rational exporters facing no current price penalty have little incentive to bear the fixed costs of switching destinations. Under this interpretation, deflection may emerge when the definitive phase introduces CBAM certificate purchases tied to the EU ETS price, which exceeded $60/tCO_2$ throughout 2024. The transitional phase serves as an information-gathering and administrative preparation period, not as a trade barrier.

Second, redirecting established bilateral trade flows involves substantial switching costs. Metal supply chains are governed by long-term contracts, quality certifications, and specialized logistics. Even if CBAM imposes a price wedge, the short-run elasticity of destination substitution may be far lower than CGE models assume. Verde (2020) notes that empirical evidence on carbon leakage through the EU ETS has generally found smaller effects than model predictions, a pattern this paper extends to carbon border adjustments.

These findings speak directly to the design of carbon border clubs (Nordhaus, 2015). If unilateral carbon border pricing generates substantial trade deflection, the case for multilateral coordination is urgent but the political incentives for free-riding are strong. If, as I find, deflection is limited even under imperfect policy instruments, unilateral action may be more effective than commonly assumed—at least in sectors with high switching costs. The policy implication is that CBAM's effectiveness should be evaluated primarily through its effect on production methods and embedded emissions, not through trade volume reallocation.

A limitation of this study is the short post-treatment window (15 months). The transitional phase was explicitly designed as a "soft start," and trade patterns may adjust more substantially as full pricing approaches. Additionally, the sample covers only metals; CBAM also applies to cement, fertilizers, and electricity, where trade patterns and switching costs differ. Future work should revisit this question as the definitive phase unfolds and as additional countries (the UK, Canada, and potentially the US) implement their own carbon border mechanisms.

6. Conclusion

The EU's Carbon Border Adjustment Mechanism represents the most significant attempt to date at unilateral climate trade policy. A central concern—voiced by both proponents of broader carbon clubs and opponents of unilateral action—is that CBAM will deflect dirty production to non-participating markets rather than reducing global emissions. This paper provides the first empirical test of this channel, using a triple-difference that exploits variation in product coverage, destination exposure, and the sharp onset of the transitional phase.

The evidence decisively rejects economically meaningful trade deflection during the first 15 months. The point estimate for the deflection margin—non-EU-bound covered imports—is 0.003 standard deviations, classified as null. This finding is robust to excluding war-affected exporters, COVID-period data, and alternative estimation methods. The result challenges the strong predictions of CGE models and suggests that, at least in the metals sector, the short-run elasticity of destination substitution is lower than theoretical models assume.

Whether this null persists as CBAM moves from reporting to pricing in 2026 remains an open and economically important question. The transitional phase provides a lower bound on the policy's trade effects; the definitive phase, with actual carbon costs tied to the EU ETS price, will provide a sharper test of the deflection hypothesis. This paper establishes the empirical baseline against which those future effects can be measured.

References

Böhringer, Christoph, Carolyn Fischer, and Knut Einar Rosendahl, "The Global Effects of Subglobal Climate Policies," *The B.E. Journal of Economic Analysis & Policy*, 2010, *10* (2).

_ , Knut Einar Rosendahl, and Halvor Briseid Storrøsten, "Robust Policies to Mitigate Carbon Leakage," *Journal of Public Economics*, 2014, *149*, 35–46.

Branger, Frédéric and Philippe Quirion, "Would Border Carbon Adjustments Prevent Carbon Leakage and Heavy Industry Competitiveness Losses? Insights from a Meta-Analysis of Recent Economic Studies," *Ecological Economics*, 2014, *99*, 29–39.

European Parliament and Council, "Regulation (EU) 2023/956 Establishing a Carbon Border Adjustment Mechanism," *Official Journal of the European Union*, 2023, *L 130/52.*

Nordhaus, William, "Climate Clubs: Overcoming Free-Riding in International Climate Policy," *American Economic Review*, 2015, *105* (4), 1339–1370.

Silva, João M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Verde, Stefano F., "The Impact of the EU Emissions Trading System on Competitiveness and Carbon Leakage: The Econometric Evidence," *Journal of Economic Surveys*, 2020, *34* (2), 320–343.

World Steel Association, *Steel Statistical Yearbook 2022*, Brussels: World Steel Association, 2022.

Table 1: Summary Statistics: Monthly Bilateral Trade Flows

Group	Period	Mean (M$)	SD (M$)	Share Zero	N
EU, Covered	Post-CBAM	247.5	226.4	0.000	210
EU, Covered	Pre-CBAM	390.8	409.8	0.000	630
EU, Uncovered	Post-CBAM	390.6	599.3	0.000	105
EU, Uncovered	Pre-CBAM	511.6	765.1	0.000	315
Non-EU, Covered	Post-CBAM	57.8	83.8	0.000	572
Non-EU, Covered	Pre-CBAM	59.8	83.9	0.000	1,855
Non-EU, Uncovered	Post-CBAM	160.1	284.3	0.000	297
Non-EU, Uncovered	Pre-CBAM	158.5	305.7	0.000	925

Notes: Monthly bilateral trade values in millions of USD from UN Comtrade. “Covered” products are iron and steel (HS 72) and aluminium (HS 76), which fall under EU CBAM. “Uncovered” products are articles of iron or steel (HS 73), which are not subject to CBAM. EU destinations comprise EU-27 member states. Non-EU destinations are the US, Japan, and the UK. Pre-CBAM: January 2020–September 2023. Post-CBAM: October 2023–December 2024. Exporters: China, Turkey, Russia, Ukraine, Vietnam, Taiwan, Brazil.

Table 2: CBAM and Trade Deflection: Triple-Difference Estimates

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) PPML
Post × EU × Covered	0.364	0.399	0.397	0.396	−0.005
	(0.379)	(0.389)	(0.397)	(0.405)	(0.095)
Cell FE (Exp×Dest×Prod)	Yes	Yes	Yes	Yes	Yes
Month FE	Yes				
Exporter×Month FE		Yes	Yes	Yes	Yes
Destination×Month FE			Yes	Yes	Yes
Product×Month FE				Yes	Yes
Observations	4,909	4,909	4,909	4,909	4,909
R^2	0.883	0.923	0.928	0.931	—

Notes: Triple-difference estimates. The dependent variable is log(trade value + 1) in columns 1–4 and trade value in levels in column 5 (PPML). "Post" = October 2023 onward (CBAM transitional phase). "EU" = EU-27 destination. "Covered" = HS 72 (iron/steel) or HS 76 (aluminium). The omitted category is uncovered products (HS 73: articles of iron/steel) shipped to non-EU destinations (US, Japan, UK). Standard errors clustered at the exporter×destination level in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Table 3: Event Study: Monthly Triple-Difference Coefficients

Event Month	Coefficient	SE	95% CI Lower	95% CI Upper
$t = -12$	-0.911	(0.567)	-2.022	0.199
$t = -9$	-1.010^{**}	(0.484)	-1.959	-0.061
$t = -6$	-0.573^{**}	(0.231)	-1.025	-0.121
$t = -3$	-0.399	(0.320)	-1.026	0.228
$t = +0$	-0.060	(0.300)	-0.648	0.528
$t = +1$	-0.017	(0.208)	-0.425	0.391
$t = +3$	-0.313	(0.292)	-0.886	0.260
$t = +6$	-0.561^{**}	(0.260)	-1.071	-0.051
$t = +9$	-0.200	(0.318)	-0.823	0.424
$t = +12$	-0.517^{**}	(0.242)	-0.992	-0.043
Pre-trend joint F-test p-value: 0.041				

Notes: Coefficients from an event study interacting monthly dummies (relative to $t = -1$, September 2023) with EU destination and CBAM-covered product indicators. The dependent variable is log(trade value +1). All models include exporter×destination×product, exporter×month, destination×month, and product×month fixed effects. Standard errors clustered at the exporter×destination level. ***, **, * denote significance at 1%, 5%, 10%. $t = 0$ corresponds to October 2023 (CBAM transitional phase onset).

Table 4: Robustness Checks

	(1) Baseline	(2) Excl. RU/UA	(3) Excl. 2020	(4) Excl. US	(5) Steel only
Post × EU × Covered	0.396	−0.065	0.288	0.388	0.281
	(0.405)	(0.185)	(0.360)	(0.451)	(0.510)
Placebo (uncovered only): Post × EU = -0.362^{*} (SE: 0.206)					
Placebo date (Oct 2022): 0.546 (SE: 0.395)					
Dose-response (CO_2 intensity): 0.413 (SE: 0.421)					
Observations	4,909	3,599	3,909	3,675	3,251
All FEs	Yes	Yes	Yes	Yes	Yes

Notes: All specifications include exporter×destination×product, exporter×month, destination×month, and product×month fixed effects. Dependent variable: log(trade value +1). Column 1 reproduces the baseline (Table 2, column 4). Column 2 excludes Russia and Ukraine (war disruption). Column 3 excludes 2020 (COVID). Column 4 excludes the US (Section 232 tariffs). Column 5 restricts to iron/steel products only (HS 72 vs HS 73). Placebo: uncovered products (HS 73) tested for differential EU trends. Placebo date: false treatment at October 2022 using only pre-CBAM data. Dose-response: interaction of the triple-difference with exporter CO_2 intensity (tCO_2/t steel). Standard errors clustered at the exporter×destination level. ***, **, * denote significance at 1%, 5%, 10%.

Appendix: Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
CBAM Deflection (DDD)	0.396	0.405	2.632	0.150	0.154	Large positive
EU-Bound Covered Imports (DiD)	-0.406	0.354	1.808	-0.225	0.196	Large negative
Non-EU-Bound Covered Imports (DiD)	0.006	0.158	2.227	0.003	0.071	Null

Notes: **Country:** European Union (EU-27) and non-EU destinations (United States, Japan, United Kingdom). **Research question:** Does the EU Carbon Border Adjustment Mechanism (CBAM) deflect high-carbon metal exports from major producers away from the EU toward non-CBAM markets? **Policy mechanism:** CBAM imposes carbon-cost reporting requirements on imports of iron/steel (HS 72) and aluminium (HS 76) entering the EU, effective October 2023 in its transitional phase; importers must report embedded emissions, creating compliance costs and signaling future carbon pricing that may redirect trade flows to unregulated destinations. **Outcome definition:** Log monthly bilateral trade value (USD) plus one, measuring the intensive margin of metal exports from major producers to EU versus non-EU destinations. **Treatment:** Binary; onset of CBAM transitional phase in October 2023. **Data:** UN Comtrade monthly bilateral imports, 7 exporters × 4 destinations × 3 product groups, January 2020–December 2024 (4,909 observations). **Method:** Triple-difference (destination × product × time) estimated via OLS with exporter×destination×product, exporter×month, destination×month, and product×month fixed effects; standard errors clustered at the exporter×destination level. **Sample:** Restricted to 7 major metal-exporting countries (China, Turkey, Russia, Ukraine, Vietnam, Taiwan, Brazil) and 4 destination markets; HS 73 (articles of iron/steel) serves as the within-sector uncovered placebo. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

Acknowledgements

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

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

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