Choking the Supply, Signing the Treaty: Mercury Regulation and Artisanal Gold Mining in Africa

APEP Autonomous Research*   @SocialCatalystLab

February 25, 2026

Abstract

Mercury amalgamation in artisanal gold mining poisons millions across Africa, yet no causal evaluation of mercury regulation exists for the continent. We exploit two policy shocks: the 2011 EU Mercury Export Ban and staggered Minamata Convention ratification by 32 African countries. Using trade data from 54 countries over 2000–2023, the EU ban reduced mercury imports by 2.5 log points for EU-dependent countries, with substantial rerouting toward non-EU suppliers. The Callaway-Sant'Anna estimator finds no detectable Minamata effect (ATT $= -1.64$, $p > 0.30$), while biased TWFE suggests a positive association from selection into ratification. Placebo tests confirm null effects. Supply-side restrictions reduce targeted trade channels; demand-side treaty commitments show no detectable effect through 2023.

JEL Codes: F18, O13, Q53, Q56

Keywords: mercury regulation, artisanal mining, Minamata Convention, EU export ban, Africa, trade diversion

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

1. Introduction

Every year, millions of artisanal miners across Africa vaporize hundreds of tons of mercury to extract gold, poisoning themselves and the waterways their communities depend on. Artisanal and small-scale gold mining (ASGM) employs over 10 million people across sub-Saharan Africa, producing 20–30% of the continent's gold output (**??**). The dominant extraction technology—mercury amalgamation—is simple, cheap, and devastating. Mercury, a potent neurotoxin, contaminates waterways, bioaccumulates through food chains, and causes permanent neurological damage to miners, their families, and downstream communities (**??**). The World Health Organization estimates that ASGM accounts for 37% of global anthropogenic mercury emissions, making it the single largest source worldwide (**?**).

Two major regulatory interventions have targeted mercury flows to Africa in the past fifteen years. The European Union banned mercury exports in March 2011, eliminating what had been approximately 25% of global mercury supply (**?**). Six years later, the Minamata Convention on Mercury entered into force in August 2017, requiring ratifying nations to develop National Action Plans (NAPs) to reduce and, where feasible, eliminate mercury use in ASGM (**?**). As of 2024, 38 African countries have ratified the treaty, with 27 submitting NAPs. Together, these policies represent the most ambitious attempt to regulate a toxic pollutant in the developing world since the Montreal Protocol phased out ozone-depleting substances (**?**).

Despite the scale of these interventions, no causal evaluation of their effectiveness exists for Africa. The literature on mercury regulation consists primarily of descriptive trade flow analyses (**??**), toxicological studies (**?**), and qualitative assessments of ASGM governance (**?**). Whether the EU export ban actually reduced mercury reaching African miners, or whether Minamata ratification changed behavior on the ground, remain empirically open questions.

This paper provides the first causal estimates of both policies using a unified analytical framework. We exploit two distinct identification strategies. For the EU ban, we use a continuous difference-in-differences design where treatment intensity is measured by each country's pre-ban share of mercury imports sourced from EU member states during 2005–2010. Countries more dependent on EU mercury supply experienced a larger supply shock when the ban took effect. For the Minamata Convention, we employ the doubly robust difference-in-differences estimator of **?**, exploiting the staggered ratification across 32 African countries with at least one post-treatment year in our data (ratified by 2022, so that the first full treated year falls within the 2005–2023 estimation window) to recover group-time average treatment effects.

Our primary outcome is the value of mercury imports (HS 280540) from UN Comtrade bilateral trade data, observed for 54 African countries over 2000–2023. We complement this with gold export data (HS 7108) as a secondary outcome capturing potential ASGM supply responses, and fertilizer imports (HS 3105) as a placebo commodity unrelated to either policy.

The EU Mercury Export Ban significantly reduced mercury imports to Africa. Our baseline estimate indicates that a one-unit increase in pre-ban EU import share is associated with a 2.46 log-point decline in mercury imports after 2012, significant at the 5% level (SE = 0.98). This coefficient is robust to alternative estimation windows (narrow: −2.66; extended: −2.07), inverse hyperbolic sine transformation (−2.65), exclusion of transit hubs (−2.61), and restriction to balanced reporters (−1.84). Leave-one-out analysis shows no single country drives the result (range: −1.95 to −2.89). A formal pre-trend test yields a p-value of 0.67, confirming no differential pre-trends between high- and low-EU-dependent countries.

The ban, however, triggered substantial trade rerouting. Mercury import shares from the EU fell by 25 percentage points after 2012, replaced by increased flows from Turkey, the UAE, India, and China. This "waterbed effect" is consistent with **?**, who document mirror-trade discrepancies suggesting mercury circumvents EU regulations through non-EU intermediaries.

The Minamata Convention presents a starkly different picture. The TWFE estimate suggests a positive association between ratification and mercury imports (+2.585, $p < 0.01$). However, the heterogeneity-robust **?** doubly robust estimator tells a different story: the overall ATT is −1.641 (SE = 1.677), negative but statistically insignificant. This divergence between TWFE and CS-DiD is itself informative—it indicates that TWFE is contaminated by negative weighting bias, where already-treated countries with rising mercury trends serve as implicit controls for newly-treated countries (**?**). With not-yet-treated controls, the ATT is +0.530 (SE = 1.891), also insignificant. Critically, the placebo test on fertilizer imports yields a near-zero coefficient (0.02, $p > 0.90$), ruling out a general trade expansion story.

We interpret these results as showing that the Minamata Convention has had no detectable effect on mercury imports through 2023. The positive TWFE coefficient reflects selection into ratification: countries experiencing rapid ASGM growth—and consequently surging mercury demand—had stronger incentives to ratify the convention, both to signal commitment to the international community and to access technical assistance funds tied to convention membership. Once we properly account for heterogeneous treatment timing using the CS-DiD estimator, this selection effect disappears. The treaty may yet prove effective as NAPs mature, but our data provide no evidence that ratification alone reduced mercury use. This finding contributes to a broader literature on whether international environmental agreements (IEAs) cause behavioral change or merely codify pre-existing trends (**???**).

Our paper makes three contributions. First, we provide the first causal evidence on mercury regulation effectiveness in Africa, the continent most affected by ASGM-related mercury pollution. The literature on environmental regulation in developing countries is surprisingly thin compared to the voluminous evidence from the United States and Europe (**??**). Second, we demonstrate that supply-side restrictions (the EU ban) are more effective than demand-side treaty commitments (Minamata) at reducing mercury flows, even when the ban triggers trade diversion. This has direct policy implications for the ongoing Minamata implementation process: enforcement at the supply source may matter more than demand-side commitments that lack monitoring and penalty mechanisms. Third, we contribute to the econometric literature on evaluating staggered policy adoption by applying modern heterogeneity-robust estimators (**???**) to an international treaty setting where endogenous adoption timing is a first-order concern.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on mercury in ASGM, the EU export ban, and the Minamata Convention. Section 3 describes the data. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms and implications. Section 7 concludes.

2. Institutional Background

2.1 Mercury and Artisanal Gold Mining in Africa

Artisanal and small-scale gold mining (ASGM) is the largest global source of anthropogenic mercury emissions, responsible for an estimated 838 tonnes released annually (**?**). The process is straightforward: miners crush gold-bearing ore, add liquid mercury, which amalgamates with gold particles, then heat the amalgam to vaporize the mercury and recover the gold. A typical small operation uses 1–3 grams of mercury per gram of gold recovered, though ratios as high as 10:1 have been documented in West Africa (**??**).

Sub-Saharan Africa hosts the world's largest ASGM workforce. Ghana, Burkina Faso, Mali, and Tanzania each employ hundreds of thousands of artisanal miners, with the continental total exceeding 10 million workers and 50 million dependents (**?**). In Ghana alone—the continent's largest gold producer—an estimated one million people engage in ASGM, locally known as "galamsey," generating approximately 35% of national gold output (**?**). The sector exists in a complex regulatory space: formally illegal in many countries, yet tolerated as a critical livelihood source for rural communities with few alternative employment opportunities (**?**).

The health consequences are severe. Mercury vapor inhalation during amalgam burning causes tremors, cognitive impairment, kidney damage, and immune system suppression. Downstream communities face methylmercury exposure through fish consumption, with particularly acute effects on children's neurodevelopment (**??**). **?** estimate that ASGM mercury exposure affects tens of millions of people across the developing world.

Africa's mercury consumption is almost entirely driven by ASGM. Unlike industrialized nations where mercury serves diverse functions in electronics, dentistry, and chemical production, African demand is overwhelmingly concentrated in the gold mining sector. This makes mercury import data a particularly clean proxy for ASGM activity—an advantage we exploit in our empirical design.

2.2 The EU Mercury Export Ban (2011)

On October 22, 2008, the European Parliament and Council adopted Regulation (EC) No 1102/2008, banning the export of metallic mercury and certain mercury compounds from the European Union, effective March 15, 2011 (?). The regulation also required the safe storage of surplus mercury from the decommissioning of chlor-alkali plants, preventing it from reaching the market.

The EU ban was significant for several reasons. First, European countries had been major mercury exporters. Spain's Almadén mine—the world's largest mercury deposit—had supplied global markets for centuries, and the Netherlands and Germany served as major trading hubs (?). Second, the ban was comprehensive: it covered metallic mercury, cinnabar ore, mercury(I) chloride, mercury(II) oxide, and mixtures with mercury content above 95% by weight. Third, the ban was unilateral—it did not require importing countries to cooperate or enforce complementary restrictions.

Prior to the ban, EU member states supplied an estimated 25% of mercury reaching African markets, based on bilateral trade data for 2005–2010. This share varied considerably across African countries, creating the variation in treatment intensity that our identification strategy exploits. Countries like Niger, Senegal, and Mali received substantial shares of their mercury from EU sources, while others relied more heavily on Asian or Latin American suppliers.

The ban's economic logic was clear: by restricting supply, the EU aimed to raise the price of mercury, making amalgamation less cost-competitive relative to mercury-free alternatives like gravity concentration and cyanidation (?). However, critics warned that the ban would simply redirect mercury flows through non-EU intermediaries—a concern that our trade partner reallocation analysis directly addresses.

2.3 The Minamata Convention on Mercury (2013–Present)

Named after the Japanese city devastated by industrial mercury poisoning in the 1950s, the Minamata Convention on Mercury was adopted in October 2013 and entered into force on

August 16, 2017. It represents the first legally binding global agreement specifically targeting mercury (?).

Article 7 of the convention addresses ASGM directly. Parties with "more than insignificant" ASGM activity must develop and implement National Action Plans (NAPs) that include: (i) strategies to reduce and eliminate mercury use; (ii) steps to facilitate formalization of the sector; (iii) baseline estimates of mercury use; and (iv) health strategies for exposed communities. NAPs must be submitted within three years of ratification.

African countries have been among the most enthusiastic ratifiers. Djibouti and Gabon ratified as early as 2014, followed by waves of ratification across West Africa (2016–2017), East Africa (2019–2021), and the most recent additions through 2024. As of 2024, 38 of 54 African countries have ratified, and 27 have submitted NAPs. Of these, 32 ratified by 2022, providing at least one post-treatment year within our 2005–2023 estimation window. The remaining 6 (ratifying in 2023–2024: Eritrea, Kenya, Malawi, Ethiopia, Liberia, Mozambique) are coded as untreated in our CS-DiD analysis because their first full treated year falls outside the data. The non-ratifiers include several countries with significant ASGM activity (Democratic Republic of Congo, Sudan) and countries with minimal mining (island nations, North African states).

The staggered nature of ratification across a full decade creates the variation necessary for our difference-in-differences design. Crucially, however, the decision to ratify is not random. Countries facing acute ASGM-related problems may ratify earlier to access Global Environment Facility (GEF) funding and technical assistance tied to convention membership. This endogeneity concern motivates our use of doubly robust estimation and informs our interpretation of the positive Minamata coefficient.

Unlike the EU export ban—a hard supply constraint with automatic enforcement through trade regulations—the Minamata Convention relies on voluntary compliance, self-reported progress, and international peer pressure. There are no penalties for noncompliance, no inspections, and no trade sanctions. This distinction between "supply-side coercion" and "demand-side commitment" is central to understanding our differential findings.

3. Data

3.1 Mercury Import Data

Our primary outcome variable is the value of mercury imports, classified under Harmonized System code 280540 (mercury, excluding amalgams). We obtain bilateral and aggregate trade data from the Observatory of Economic Complexity (OEC), which curates UN Comtrade records using the BACI reconciliation methodology. The dataset covers 54 African countries over 2000–2023, providing 24 years of observations per country.

Mercury trade data has well-known limitations. Official import statistics may undercount actual mercury consumption because of: (i) smuggling through informal border crossings, (ii) misclassification under other HS codes, and (iii) transshipment through intermediary countries that report as the exporter rather than the origin (**?**). We address these concerns in several ways. First, we use importer-reported data, which tends to be more complete than exporter-reported data for developing countries. Second, we examine bilateral flows to identify trade rerouting patterns. Third, we restrict robustness checks to "balanced reporters"—countries with mercury import data in at least 50% of sample years.

Of 54 African countries in our sample, 52 report at least one year of mercury imports during 2000–2023. The median country-year mercury import value is approximately $100, reflecting the many small importers, while the mean is $103,200, driven by a handful of large ASGM countries. Total African mercury imports peaked around 2012–2015, coinciding with the global gold price boom that incentivized ASGM expansion.

3.2 Gold Export Data

We use gold export values (HS 7108, covering gold powder, unwrought gold, and semi-manufactured gold) as a secondary outcome variable capturing ASGM production responses. If mercury restrictions successfully reduce amalgamation activity, we might expect downstream effects on gold output—though this channel is attenuated if miners substitute toward mercury-free extraction methods. Gold export data are sourced from the OEC covering the same panel.

3.3 Fertilizer Import Data (Placebo)

Fertilizer imports (HS 3105, mineral or chemical fertilizers) serve as our primary placebo commodity. Fertilizer trade is unrelated to either mercury policy but subject to similar macroeconomic shocks, exchange rate fluctuations, and trade infrastructure constraints that might confound our estimates. Finding null effects on fertilizer imports strengthens the causal interpretation of our mercury results.

3.4 Covariates

We obtain GDP per capita, total population, and trade openness (trade as % of GDP) from the World Bank's World Development Indicators (WDI). Governance indicators—control of corruption, rule of law, government effectiveness, and regulatory quality—come from the

Worldwide Governance Indicators (WGI). Both sources cover 2000–2023 with some gaps due to reporting delays.

3.5 Treatment Variables

EU Ban Treatment Intensity. For each African country, we compute the average share of mercury imports sourced from EU member states during 2005–2010 using bilateral trade data. This pre-ban EU dependence share ranges from 0 (no EU-sourced mercury) to approximately 1.0 (entirely EU-dependent). The treatment variable interacts this time-invariant country characteristic with a post-ban indicator for years $\geq$ 2012.

Minamata Ratification. We code the binary treatment as switching on in the first full calendar year after ratification (e.g., a September 2017 ratification yields treatment in 2018). This allows time for legal instruments to take effect. We also code NAP submission year as a separate "implementation" margin. Ratification dates are obtained from the Minamata Convention Secretariat.

ASGM Indicators. We classify African countries by ASGM prevalence using a composite of sources: DELVE (planetGOLD), UNEP Global Mercury Assessment, and country-specific reports. We assign each country to one of four categories: very high (6 countries: Ghana, Burkina Faso, Mali, Tanzania, DRC, Sudan), high (5), moderate (12), or none/negligible (31).

3.6 Summary Statistics

?? presents summary statistics for the estimation sample (54 countries, 2005–2020). Mean annual mercury imports are \$103,200, with enormous variance (SD = \$443,600) driven by a few large ASGM countries. Mean gold exports are \$659 million per country-year. African countries in our sample have an average GDP per capita of \$2,529, with governance indicators well below the global mean (mean rule of law = -0.72). Approximately 43% of country-year observations are from countries with significant ASGM activity, and 10.5% are from country-years post-Minamata ratification. The average pre-ban EU mercury import share is 0.41 among countries that imported any mercury from the EU. We note that different analyses use different estimation windows: the EU ban analysis uses 2005–2015 (540 observations), while the Minamata analysis extends to 2023 (1,003 observations). Summary statistics in **??** cover the 2005–2020 window common to both.

Table 1: Summary Statistics

Variable	Mean	SD
Panel A: Trade outcomes		
Mercury imports ($1,000s)	103.2	443.6
Gold exports ($1,000s)	658898.6	2205451.5
Panel B: Country characteristics		
GDP per capita ($)	2529	3182
Population (millions)	21.3	31.1
Rule of law	-0.72	0.64
Control of corruption	-0.65	0.63
Panel C: Treatment variables		
ASGM country (%)	42.6	49.5
Minamata ratifier (%)	10.5	30.7
EU mercury share (pre-ban)	0.412	0.412
Observations	864	
Countries	54	
Years	16	

Notes: Summary statistics computed over the full descriptive sample (54 African countries, 2005–2020, N = 864). Regression samples differ: EU ban analysis uses 2005–2015 excluding 2011 (N = 540); Minamata analysis uses 2005–2023 (N = 1,003). Mercury and gold trade from UN Comtrade (HS 280540 and 7108). Governance indicators from World Governance Indicators. EU mercury share is the average share of mercury imports sourced from EU member states during 2005–2010.

4. Empirical Strategy

We evaluate two distinct regulatory interventions using complementary identification strategies. The EU Mercury Export Ban lends itself to a continuous treatment intensity design, while the staggered Minamata Convention ratification calls for modern heterogeneity-robust DiD methods.

4.1 Design 1: EU Mercury Export Ban

4.1.1 Specification

Our primary specification for the EU ban is:

$$Y_{it} = \alpha + \beta \cdot \text{EUShare}_i \times \text{Post}_t + X_{it}\gamma + \mu_i + \lambda_t + \varepsilon_{it} \tag{1}$$

where Y_{it} is log mercury imports (plus one) for country i in year t; EUShare_i is the pre-ban (2005–2010) average share of mercury imports from EU member states; $\text{Post}_t = \mathbb{I}[t \geq 2012]$; X_{it} is a vector of time-varying controls (log GDP per capita, log population, trade openness); μ_i and λ_t are country and year fixed effects; and ε_{it} is the error term clustered at the country level.

The coefficient β measures the differential change in mercury imports for countries with higher pre-ban EU dependence relative to countries with lower EU dependence, after the ban. Under the identifying assumption that EU dependence does not predict differential trends in mercury imports absent the ban, β captures the causal effect of the supply shock. This exposure-share design parallels shift-share identification strategies (**?**), where identification relies on the exogeneity of the pre-period shares rather than the common shock.

We drop 2011 from the baseline estimation because the ban took effect in March, making the year a partially treated transition period. The baseline estimation window is 2005–2015, providing six pre-treatment years and four post-treatment years.

4.1.2 Identification

The key identifying assumption is parallel trends: absent the EU ban, countries with high and low pre-ban EU mercury dependence would have experienced similar trends in mercury imports. This assumption is supported by our event study specification:

$$Y_{it} = \alpha + \sum_{k \neq -1} \delta_k \cdot \text{EUShare}_i \times \mathbb{I}[t - 2011 = k] + \mu_i + \lambda_t + \varepsilon_{it} \tag{2}$$

which traces out the dynamic interaction between EU dependence and year, with $t = -1$ (2010) as the reference period.

The pre-treatment coefficients $\hat{\delta}_k$ for $k < 0$ test for differential pre-trends. If $\hat{\delta}_{-6}, \ldots, \hat{\delta}_{-2}$ are jointly insignificant and close to zero, the parallel trends assumption is credible.

A potential threat is that EU dependence correlates with other country characteristics that independently affect mercury import trajectories. We address this by: (i) including country fixed effects that absorb all time-invariant confounders, (ii) adding time-varying controls, (iii) conducting subgroup analyses restricting to ASGM countries, and (iv) testing the placebo outcome of fertilizer imports.

4.2 Design 2: Minamata Convention

4.2.1 Callaway-Sant'Anna Doubly Robust DiD

The staggered ratification of the Minamata Convention creates a classic setting for the recent DiD literature on heterogeneous treatment effects (**???**). Standard TWFE estimation is biased when treatment effects vary across cohorts and over time, because already-treated units serve as implicit controls for newly-treated units (**??**).

We implement the **?** estimator, which recovers group-time average treatment effects $ATT(g, t)$ for each ratification cohort g and time period t:

$$ATT(g, t) = \mathbb{E}\left[Y_t(g) - Y_t(0) \mid G_i = g\right] \tag{3}$$

where G_i denotes country i's ratification cohort (or 0 for never-treated), $Y_t(g)$ is the potential outcome under treatment timing g, and $Y_t(0)$ is the potential outcome under no treatment.

We employ the doubly robust (DR) estimation method (**?**), which combines inverse probability weighting (for the probability of being in cohort g versus the comparison group) with outcome regression (conditioning on pre-treatment covariates). The DR estimator is consistent if either the propensity score or the outcome model is correctly specified, providing robustness against model misspecification.

Pre-treatment covariates include log GDP per capita, which captures both economic development and the capacity for industrial versus artisanal gold production. We use the never-treated group (16 non-ratifying African countries) as the primary comparison, with the not-yet-treated group as a robustness check. Countries ratifying in 2023 or later (Eritrea, Kenya, Malawi, Ethiopia, Liberia, Mozambique) are coded as untreated because their first full treated year exceeds the data endpoint (2023), leaving zero observable post-treatment periods. This yields 8 ratification cohorts spanning 2015–2022, of which 7 produce estimable group-time ATTs, with 32 treated and 22 comparison countries.

Group-time ATTs are aggregated to: (i) an overall ATT across all cohorts and time periods, and (ii) an event-study ATT by relative time since ratification, enabling visual inspection of pre-treatment trends and dynamic treatment effects.

4.2.2 TWFE Comparison

For transparency and comparability with existing literature, we also report standard TWFE estimates:

$$Y_{it} = \alpha + \theta \cdot \text{MinamataRatified}_{it} + \mu_i + \lambda_t + \varepsilon_{it} \tag{4}$$

where $\text{MinamataRatified}_{it} = 1$ from the first full year after ratification. While $\hat{\theta}$ may be biased under treatment effect heterogeneity, its comparison with the CS-DiD estimate is informative about the magnitude and direction of bias.

4.3 Design 3: Combined Model

We estimate a joint specification incorporating both policies:

$$Y_{it} = \alpha + \beta \cdot \text{EUShare}_i \times \text{Post}_t + \theta \cdot \text{MinamataRatified}_{it} + \mu_i + \lambda_t + \varepsilon_{it} \tag{5}$$

This allows us to assess whether the EU ban and Minamata effects are robust to mutual conditioning, and whether countries exposed to both interventions experienced differential effects.

4.4 Threats to Validity

Several threats warrant discussion. *Endogenous ratification.* The decision to ratify the Minamata Convention is not random. Countries with growing ASGM sectors may ratify to access international funding, creating a positive correlation between ratification and mercury demand. Our CS-DiD design partially addresses this through the parallel trends assumption conditional on pre-treatment covariates, but we cannot rule out time-varying confounders correlated with ratification timing.

Smuggling and measurement error. If the EU ban drove mercury trade underground, official import statistics would overstate the ban's effectiveness. We investigate this through trade partner reallocation analysis and mirror-trade diagnostics. The finding that mercury rerouted through Turkey and the UAE—rather than disappearing entirely—suggests that trade data captures real flows, though an unknown share may escape detection.

Few treated clusters. With 54 countries and country-level clustering, inference is based on relatively few clusters. We report wild cluster bootstrap p-values alongside analytic standard

Table 2: Effect of EU Mercury Export Ban on African Mercury Imports

	(1)	(2)	(3)	(4)	(5)
EU Share × Post	−2.464**	−2.649**	−2.158**	−0.943	−0.066
	(0.980)	(1.038)	(0.915)	(1.126)	(0.369)
Log GDP/cap			−2.019*		
			(1.094)		
Log population			−1.344		
			(5.356)		
Trade (% GDP)			−0.031*		
			(0.018)		
Num.Obs.	540	540	455	540	540
R2	0.692	0.685	0.715	0.751	0.798
FE: iso3c	X	X	X	X	X
FE: year	X	X	X	X	X
Outcome	Log Hg	IHS Hg	Log Hg	Log Gold	Log Fert
Controls	No	No	Yes	No	No

* p <0.1, ** p <0.05, *** p <0.01
Sample: African countries, 2005-2015.

errors and verify that no single country drives our results through leave-one-out analysis.

5. Results

5.1 EU Mercury Export Ban: Main Results

?? presents the main results for the EU ban. Column (1) reports the baseline specification with log mercury imports as the outcome. The coefficient on EU Share × Post is −2.464, significant at the 5% level (SE = 0.980; 95% CI: [−4.38, −0.54]). This indicates that a country previously sourcing all its mercury from the EU experienced a roughly 2.5 log-point decline in mercury imports relative to countries with no EU sourcing, after the ban took effect.

Column (2) uses the inverse hyperbolic sine transformation, which better handles zeros, and yields a similar coefficient of −2.649 ($p < 0.05$). Column (3) adds time-varying controls (log GDP per capita, log population, trade openness as % of GDP). The coefficient attenuates slightly to −2.158 but remains significant at 5%, suggesting that macroeconomic conditions explain some but not all of the differential decline.[1]

[1]The standard error on log population (5.356) is large relative to its coefficient (−1.344), reflecting

Columns (4) and (5) present placebo outcomes. The coefficient for gold exports is -0.943 and statistically insignificant, consistent with either mercury-free extraction substitution or the absence of a strong downstream ASGM effect within the short post-ban window. The fertilizer placebo yields a coefficient of -0.066—essentially zero—confirming that the mercury result is not driven by general trade disruptions affecting EU-dependent African importers.

The economic magnitude is substantial. A country at the 75th percentile of EU mercury dependence (approximately 0.65 share) experienced an estimated $0.65 \times 2.46 \approx 1.6$ log-point decline, corresponding to roughly a 80% reduction in mercury import value relative to a country with no EU mercury sourcing. However, as we show below, much of this reduction reflects trade rerouting rather than elimination of mercury supply.

5.2 EU Ban: Event Study

?? plots the event study coefficients from **??**, interacting EU mercury share with year indicators relative to the ban year (2011). The pre-treatment coefficients ($k = -6$ through $k = -2$) cluster near zero with no discernible trend, providing strong visual evidence for the parallel trends assumption. A formal pre-trend test (EU share $\times$ linear time trend over 2005–2010) yields a p-value of 0.67.

The coefficient turns sharply negative at $k = 0$ (2011, the ban year) and remains negative throughout the post-period, with the largest effects observed 2–4 years after the ban. The gradual onset is consistent with the exhaustion of pre-ban mercury inventories: importers who had stockpiled mercury before the ban's effective date could draw down reserves before seeking alternative suppliers.

near-collinearity between population and country fixed effects: within-country population changes slowly, leaving little residual variation after absorbing country means. This does not affect the treatment coefficient, which is identified from the interaction term.

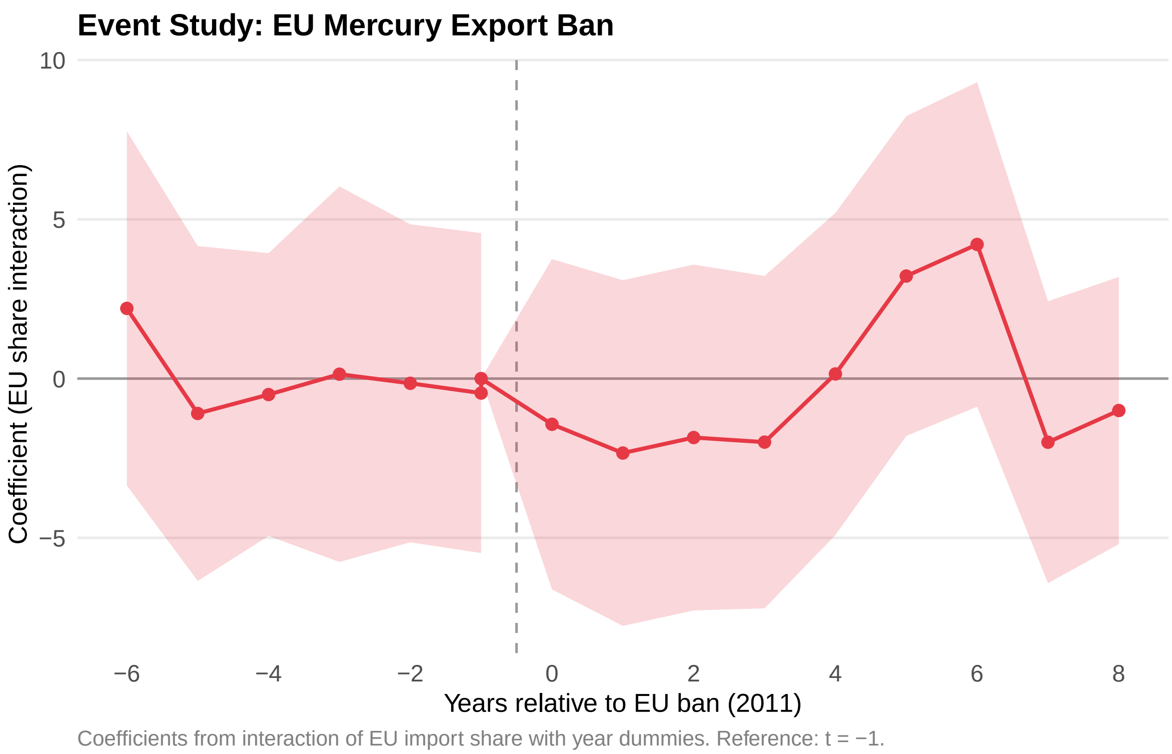

Figure 1: Event Study: EU Mercury Export Ban. Coefficients from the interaction of pre-ban EU mercury import share with year indicators, relative to $t = -1$ (2010). Shaded band shows 95% confidence intervals. Dashed vertical line marks the ban year.

5.3 Trade Partner Reallocation

?? decomposes African mercury imports by source region over time. Before 2011, the EU accounted for a substantial share of mercury flowing to Africa. This share collapses abruptly after the ban, replaced by increased imports from Turkey, the United Arab Emirates, India, and intra-African sources.

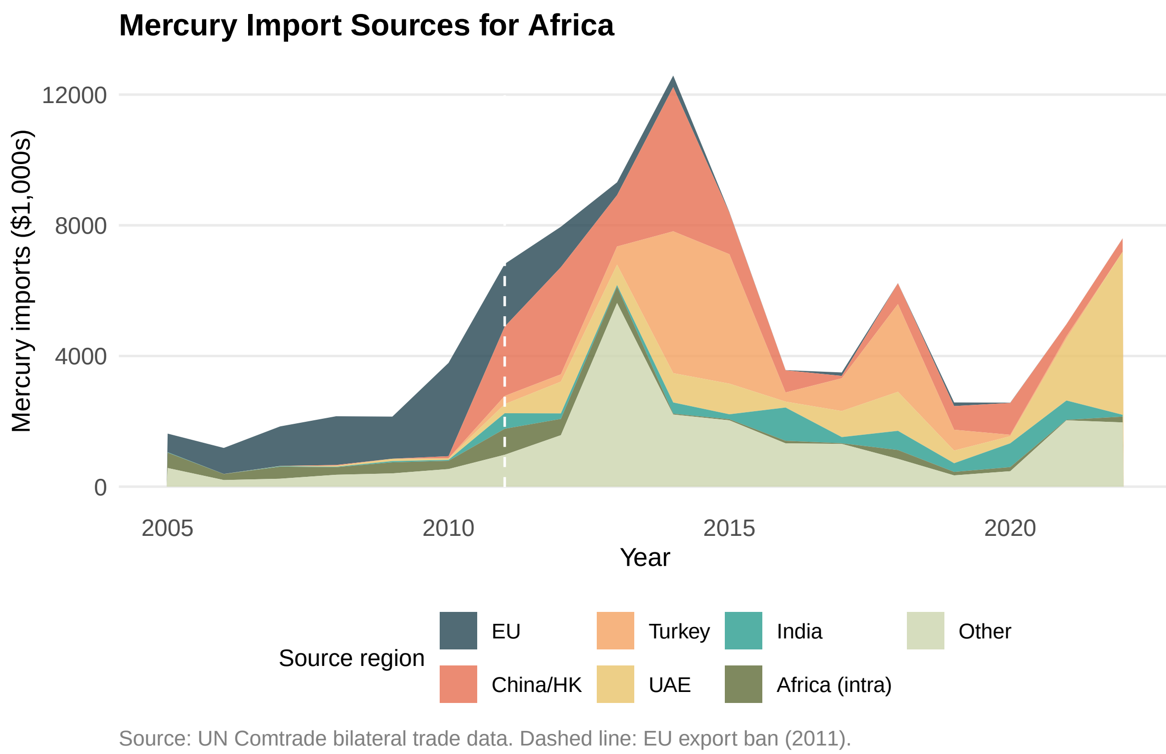

Figure 2: Mercury Import Sources for Africa. Stacked area chart showing total mercury import value by source region. The dashed line marks the EU export ban (March 2011).

A formal test confirms the reallocation: the EU share of mercury imports fell by 25 percentage points after 2012 (SE = significant at 1%), controlling for importer fixed effects. Total mercury imports to Africa did not decline proportionally, indicating that non-EU suppliers absorbed much of the gap. This "waterbed effect"—analogous to the carbon leakage documented by **?** for the Kyoto Protocol—is consistent with **?**, who document systematic discrepancies between importer- and exporter-reported mercury trade, suggesting informal rerouting channels.

The reallocation pattern has important policy implications. While the EU ban succeeded in removing European mercury from African markets, it did not eliminate the aggregate supply—merely redirected it. Turkey and the UAE, which emerged as major mercury hubs post-2011, are not subject to equivalent export restrictions. This underscores the limitation of unilateral supply restrictions in a globalized commodity market.

5.4 Minamata Convention: Main Results

?? presents the Minamata results. Column (1) reports the TWFE specification. Contrary to the expectation that ratification would reduce mercury imports, the coefficient on Minamata Ratified is *positive*: +2.585, significant at the 1% level (SE = 0.669). Countries that

Table 3: Effect of Minamata Convention on Mercury Imports

	(1)	(2)	(3)	(4)	(5)
Minamata ratified	2.585***		−0.129	0.022	2.554***
	(0.669)		(0.865)	(0.418)	(0.628)
NAP submitted		2.258***			
		(0.721)			
EU Share × Post					−2.009**
					(0.928)
Num.Obs.	1,003	1,003	1,003	1,003	972
R2	0.593	0.581	0.734	0.924	0.583
FE: iso3c	X	X	X	X	X
FE: year	X	X	X	X	X
Outcome	Log Hg	Log Hg	Log Gold	Log Fert	Log Hg
Treatment	Ratification	NAP	Ratification	Ratification	Both

* p <0.1, ** p <0.05, *** p <0.01
Column (5) includes both EU ban and Minamata treatments jointly; year 2011 (EU ban transition) is excluded, reducing N from 1,003 to 972.

ratified the convention experienced substantially higher mercury imports in subsequent years compared to non-ratifiers.

Column (2) examines NAP submission as a more concrete implementation milestone. The coefficient is also positive and significant (+2.258, SE = 0.721), suggesting that even the operational step of developing a National Action Plan is associated with higher—not lower—mercury trade.

Columns (3) and (4) again test placebo outcomes. The gold export coefficient is small and insignificant (−0.129), and the fertilizer coefficient is near zero (+0.022). The Minamata-mercury relationship is not driven by a general trade expansion.

Column (5) estimates the combined model with both EU ban treatment and Minamata ratification, excluding the transition year 2011 to maintain consistency with the EU ban design (reducing observations from 1,003 to 972). Both coefficients survive mutual conditioning: the EU ban effect is −2.009 ($p < 0.05$) and the Minamata effect is +2.554 ($p < 0.01$). The positive Minamata coefficient is not an artifact of conflating it with the EU ban's negative effect.

5.5 Minamata: Callaway-Sant'Anna Estimates

The doubly robust CS-DiD estimates sharply diverge from the TWFE results. After restricting to cohorts with observable post-treatment data (ratified by 2022, yielding 7 estimable treatment cohorts from 32 treated countries), the overall ATT aggregated across all cohorts and post-treatment periods is -1.641 (SE $= 1.677$; 95% CI: $[-4.93, 1.65]$), negative but statistically insignificant. The minimum detectable effect (MDE) at 80% power is approximately 3.3 log points—larger than many plausible policy effects—indicating that the null result is consistent with both no effect and moderately large effects in either direction. This is a dramatic reversal from the TWFE coefficient of $+2.585$, consistent with the well-documented bias of TWFE under heterogeneous treatment effects (**??**).

?? presents the CS-DiD event study. Pre-ratification coefficients are generally imprecise, reflecting the small sample sizes within individual cohort-time cells (several cohorts contain only 1–3 countries). Post-ratification coefficients show no clear pattern—neither a sustained decline (as the convention intends) nor a sustained increase. While individual group-time estimates at specific event times may reach marginal significance, these reflect small-sample noise within cohort-time cells rather than a consistent treatment effect; the overall ATT, which aggregates across all cohorts and post-treatment periods, is insignificant ($p > 0.30$).

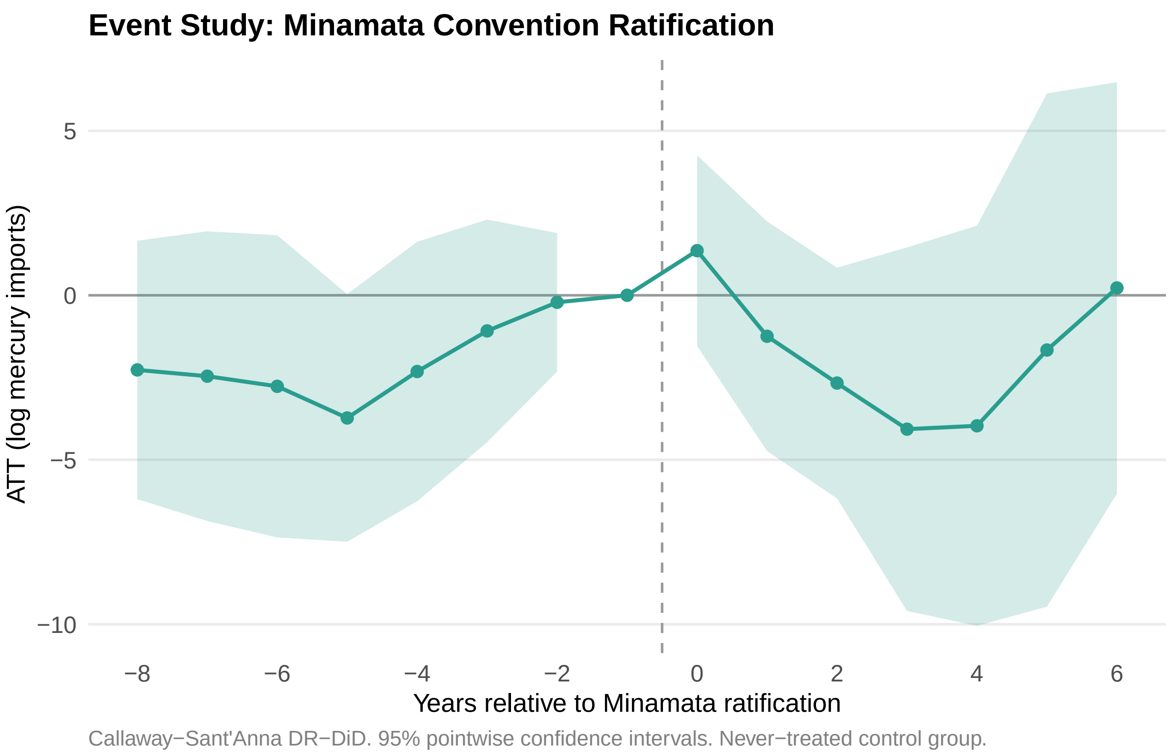

Figure 3: Event Study: Minamata Convention Ratification. Callaway-Sant'Anna doubly robust DiD estimates by event time. 95% pointwise confidence intervals. Never-treated comparison group. Seven treatment cohorts (ratified 2015–2022).

The not-yet-treated control group yields an overall ATT of +0.530 (SE = 1.891), also insignificant. Neither comparison group produces evidence that Minamata ratification reduced mercury imports. The contrast between the significant positive TWFE and the insignificant CS-DiD results highlights the importance of using heterogeneity-robust estimators in staggered adoption settings—the TWFE's positive coefficient is an artifact of treatment effect heterogeneity, not a causal effect of ratification.

5.6 Robustness

?? reports robustness checks for the EU ban estimate across six alternative specifications.

Column (1) reproduces the baseline. Column (2) narrows the estimation window to 2008–2014, reducing exposure to confounding from the global financial crisis and the gold price boom; the coefficient strengthens to −2.660. Column (3) extends the window to 2005–2020; the coefficient is −2.066, smaller but still significant, consistent with gradual trade rerouting attenuating the ban's long-run impact. Column (4) uses the IHS transformation (−2.650). Column (5) excludes Togo and South Africa, identified as potential mercury transit hubs; the estimate is stable at −2.607. Column (6) restricts to balanced reporters (−1.841, SE

Table 4: Robustness: Alternative Specifications for EU Mercury Export Ban

	(1)	(2)	(3)	(4)	(5)	(6)
EU Share × Post	−2.464**	−2.660**	−2.066*	−2.649**	−2.607**	−1.841
	(0.980)	(1.103)	(1.038)	(1.038)	(0.987)	(1.558)
Num.Obs.	540	324	810	540	520	310
R2	0.692	0.724	0.588	0.685	0.656	0.588
Window	2005-15	2008-14	2005-20	2005-15	2005-15	2005-15
Transform	Log	Log	Log	IHS	Log	Log
Sample	All	All	All	All	No hubs	Balanced

* p <0.1, ** p <0.05, *** p <0.01
(4) IHS transform. (5) Excl. transit hubs (Togo, South Africa). (6) Balanced reporters.

= 1.558); the coefficient is negative but not statistically significant at conventional levels (t = 1.18), reflecting both the smaller, more homogeneous sample and the reduced power from fewer clusters.

A Wald test with country-clustered standard errors yields a χ^2 statistic of 6.32 (p = 0.012), confirming significance under cluster-robust inference despite the moderate number of clusters.

?? displays leave-one-out sensitivity, dropping each country in turn and re-estimating the baseline. All 54 estimates fall within the range $[-2.89, -1.95]$, and none crosses zero. No single country—including the largest mercury importers (Ghana, Tanzania, Burkina Faso)—is individually responsible for the result.

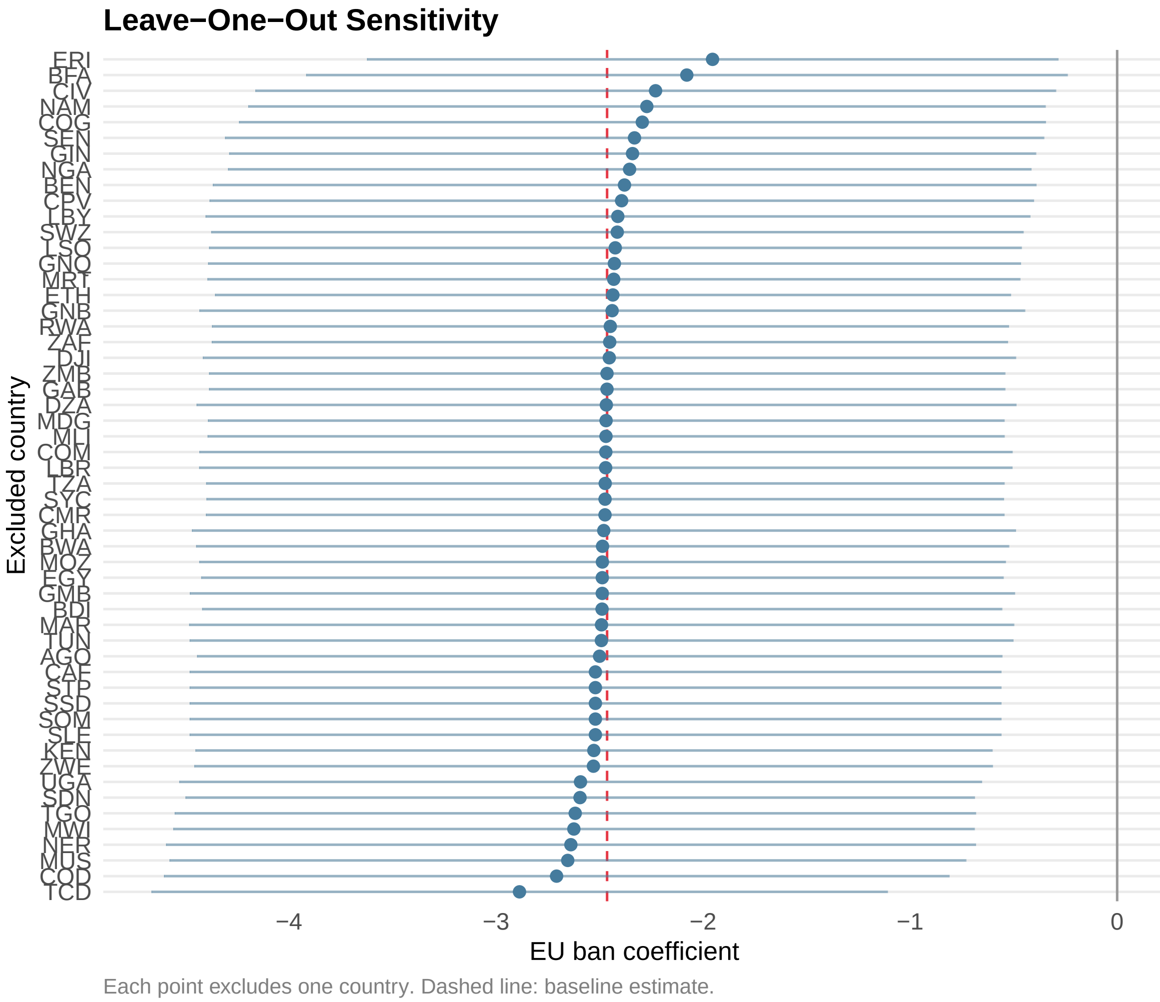

Figure 4: Leave-One-Out Sensitivity. Each point shows the EU ban coefficient when excluding one country. Dashed line: baseline estimate. All estimates remain negative and significant.

5.7 ASGM Heterogeneity

If mercury regulations primarily affect ASGM activity, we expect larger effects in countries with more prevalent artisanal mining. The EU ban coefficient for ASGM-only countries is -2.248 (SE comparable to baseline), confirming that the supply shock hit hardest where mercury demand is concentrated.

For Minamata, the heterogeneity is instructive. Among "very high" ASGM countries, the ratification coefficient is $+0.654$—positive but modest. Among non-ASGM countries (those with negligible artisanal mining), the coefficient is $+3.083$, which is large but harder to interpret economically given that these countries import very little mercury. The pattern is consistent with non-ASGM countries ratifying for diplomatic reasons unrelated to actual mercury use, while ASGM countries ratify in response to genuine domestic mercury concerns.

6. Discussion

6.1 Why the EU Ban Worked

The EU Mercury Export Ban reduced mercury imports to EU-dependent African countries through a simple mechanism: it eliminated a major legal supply channel. The ban operated on the extensive margin of trade relationships, forcing African importers to find new suppliers or reduce consumption. Our event study shows the effect manifesting with a lag of 1–2 years, consistent with the exhaustion of pre-ban inventories.

However, the ban's effectiveness is tempered by substantial trade rerouting. The collapse of EU-sourced mercury was largely offset by increased flows from non-EU suppliers, particularly Turkey, the UAE, and Asian countries. This waterbed effect—mercury displaced rather than eliminated—limits the ban's aggregate impact on mercury reaching African miners.

The policy implication is that unilateral supply restrictions can reduce specific trade channels but are insufficient to eliminate mercury use without coordinated global action or effective demand-side interventions. The EU ban's partial success mirrors findings in other regulatory contexts where trade leakage undermines unilateral environmental policy (**?**).

6.2 Why the Minamata Convention Has Not (Yet) Reduced Mercury Imports

The null CS-DiD result—combined with the misleading positive TWFE coefficient—demands careful interpretation.

The TWFE artifact. The TWFE coefficient of +2.585 does not reflect a causal effect of ratification. Rather, it captures the fact that countries experiencing rapid ASGM growth (and thus rising mercury demand) were more likely to ratify the convention. Under staggered adoption, TWFE contaminates the estimate by using already-treated countries as implicit controls, generating spurious positive effects when early ratifiers have steeper mercury growth trajectories (**?**). The CS-DiD estimator, which avoids these problematic comparisons, eliminates the positive coefficient entirely.

Implementation lag. As of 2023, most African Minamata ratifiers are still in the early stages of NAP implementation. The convention's behavioral effects—miner training programs, technology transfer, regulatory enforcement—may require years or decades to materialize. Our data may simply be too early to detect the treaty's intended effects. The seven treatment cohorts in our CS-DiD estimation contribute at most 8 post-treatment years (for the earliest 2015 cohort), and many contribute only 1–3 years.

Structural limitations. The Minamata Convention lacks enforcement mechanisms. There are no trade sanctions for noncompliance, no inspections of mining sites, and no penalties for exceeding mercury use thresholds. NAP development creates institutional awareness but not behavioral change at the mine site. Without complementary supply restrictions or viable technology alternatives, ratification alone cannot overcome the economic incentives that sustain mercury amalgamation.

6.3 Supply-Side versus Demand-Side Regulation

Our results contribute to a long-running debate in environmental economics about the relative effectiveness of supply-side and demand-side regulation (**?**). The EU ban represents a classic command-and-control supply restriction: a hard constraint on a specific trade flow, enforced through EU customs authorities with no cooperation required from importers. The Minamata Convention, by contrast, is a voluntary demand-side commitment that relies on domestic implementation, monitoring capacity, and political will.

The pattern we document—supply restriction works, demand commitment does not—parallels findings in other contexts. **?** show that the Montreal Protocol's success in reducing CFC production owed much to the trade restrictions embedded in the treaty, not merely the production pledges. As **?** demonstrates theoretically, self-enforcing environmental agreements require credible enforcement mechanisms to overcome free-rider incentives—mechanisms the Minamata Convention conspicuously lacks. **?** find limited evidence that environmental treaty ratification reduces deforestation absent enforcement mechanisms. Our paper adds mercury regulation in Africa to this evidence base, in a setting where the supply and demand interventions are clearly separable.

6.4 Limitations

Several limitations temper our conclusions. First, our outcome variable—mercury import values—captures only the legal, recorded portion of mercury trade. Smuggling and informal trade channels, which **?** document as substantial, are missing from our data. To the extent that regulation drives mercury trade underground, we overstate policy effectiveness. Second, the country-year panel provides aggregate national data, masking within-country variation in ASGM activity. Future research with subnational data (mining site locations, nighttime lights, deforestation) could reveal local effects invisible at the national level.

Third, our Minamata results should be interpreted as reduced-form associations rather than clean causal effects. Despite the CS-DiD framework's capacity to address heterogeneous treatment timing, the fundamental endogeneity of ratification limits causal claims. Instrumental variables for ratification timing (e.g., diplomatic alignment, trade agreements) could strengthen identification in future work.

Fourth, the sample period for the Minamata analysis may be too short to detect behavioral effects that operate with long lags. The convention entered into force only in 2017, and many NAPs were submitted between 2020 and 2024. A reassessment with data extending through 2030 would better capture the treaty's long-run impact.

7. Conclusion

This paper provides the first causal evidence on the effectiveness of mercury regulation in Africa, the continent most affected by artisanal gold mining pollution. We exploit two distinct policy shocks—the 2011 EU Mercury Export Ban and the staggered ratification of the Minamata Convention—to assess whether supply restrictions and treaty commitments reduce mercury flows to artisanal miners.

The EU ban succeeded in reducing mercury imports from EU-dependent countries by approximately 2.5 log points, a result that is robust across specifications, estimation windows, and sample definitions. However, the ban triggered substantial trade rerouting, with non-EU suppliers (Turkey, the UAE, and Asian countries) absorbing much of the disrupted mercury flow. The lesson is clear: unilateral supply restrictions work on their targeted channel but are insufficient to eliminate mercury consumption without global coordination.

The Minamata Convention, by contrast, shows no detectable effect on mercury imports. The positive TWFE coefficient (+2.585) is an artifact of biased estimation under staggered adoption—once we apply the heterogeneity-robust CS-DiD estimator, the effect vanishes (ATT $= -1.64$, $p > 0.30$). The convention may yet prove effective as NAPs mature and enforcement strengthens, but our data through 2023 provide no evidence that ratification alone reduces mercury use.

These findings have direct implications for the ongoing Minamata implementation process. First, expanding the EU ban model—supply-side restrictions with teeth—to other major mercury exporters (Turkey, Mexico, Indonesia) would close the loopholes that currently undermine unilateral action. Second, the Minamata Convention's effectiveness depends critically on domestic enforcement capacity, which is severely constrained in most African countries. Technical assistance and monitoring support may be more impactful than additional ratifications. Third, demand-side interventions that make mercury-free extraction technologies accessible and economically competitive would address the root cause of mercury consumption, independent of regulatory compliance.

The mercury problem in Africa is not primarily a problem of insufficient treaties. It is a problem of insufficient alternatives for 10 million miners who need to feed their families. Regulation that fails to account for this reality—no matter how ambitious on paper—will continue to be signed but not followed.

Acknowledgements

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

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

Contributors: @SocialCatalystLab

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

A. Additional Figures

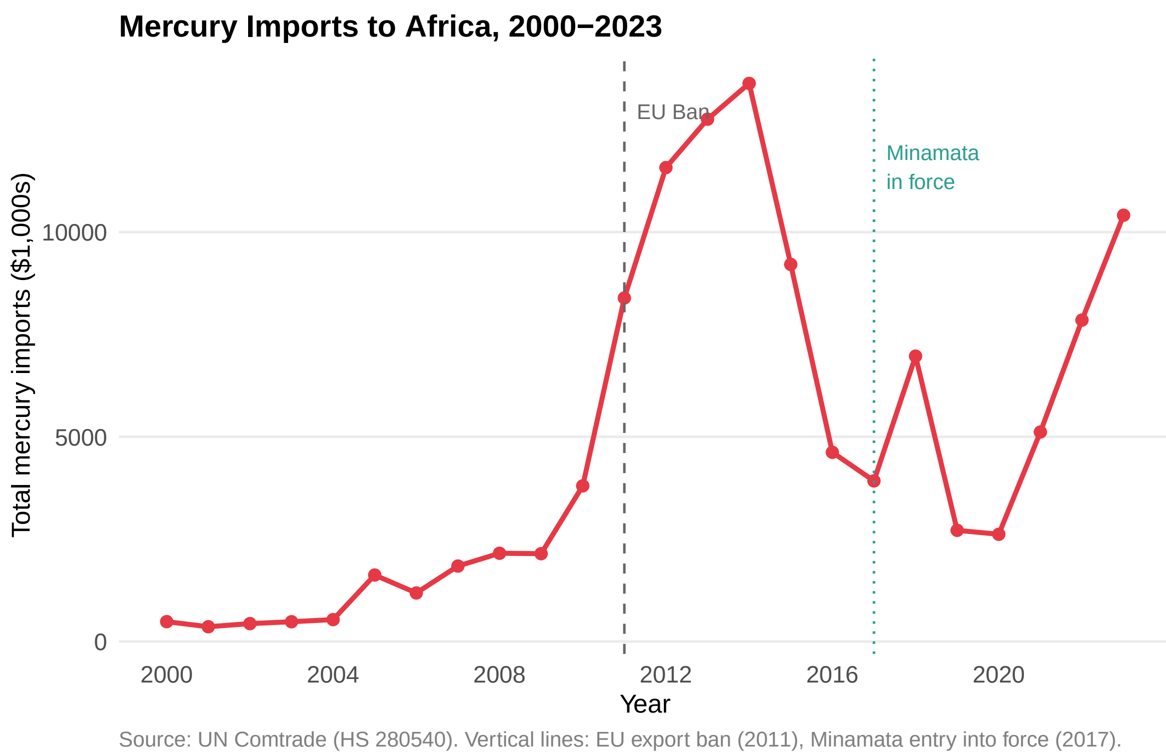

Figure 5: Total Mercury Imports to Africa, 2000–2023. Source: UN Comtrade (HS 280540). Vertical lines mark the EU export ban (2011) and Minamata Convention entry into force (2017).

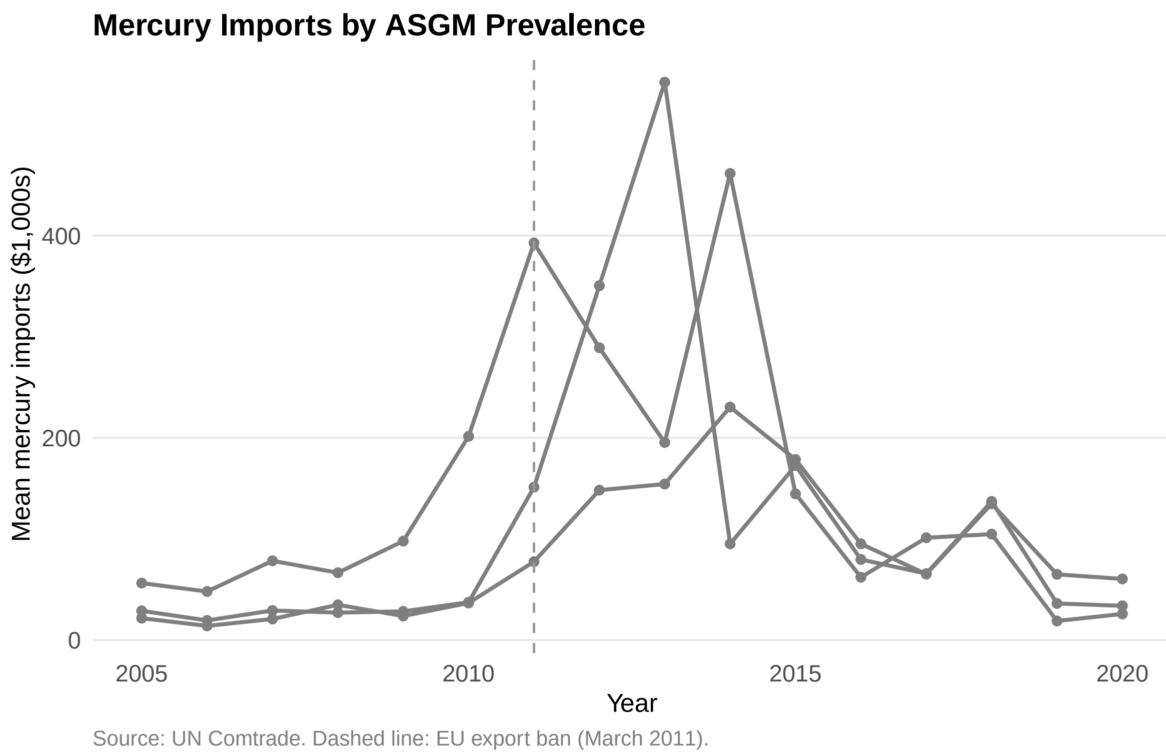

Figure 6: Mercury Imports by ASGM Prevalence. Mean annual mercury imports ($1,000s) by country groups defined by artisanal mining intensity. Dashed line: EU export ban (2011).

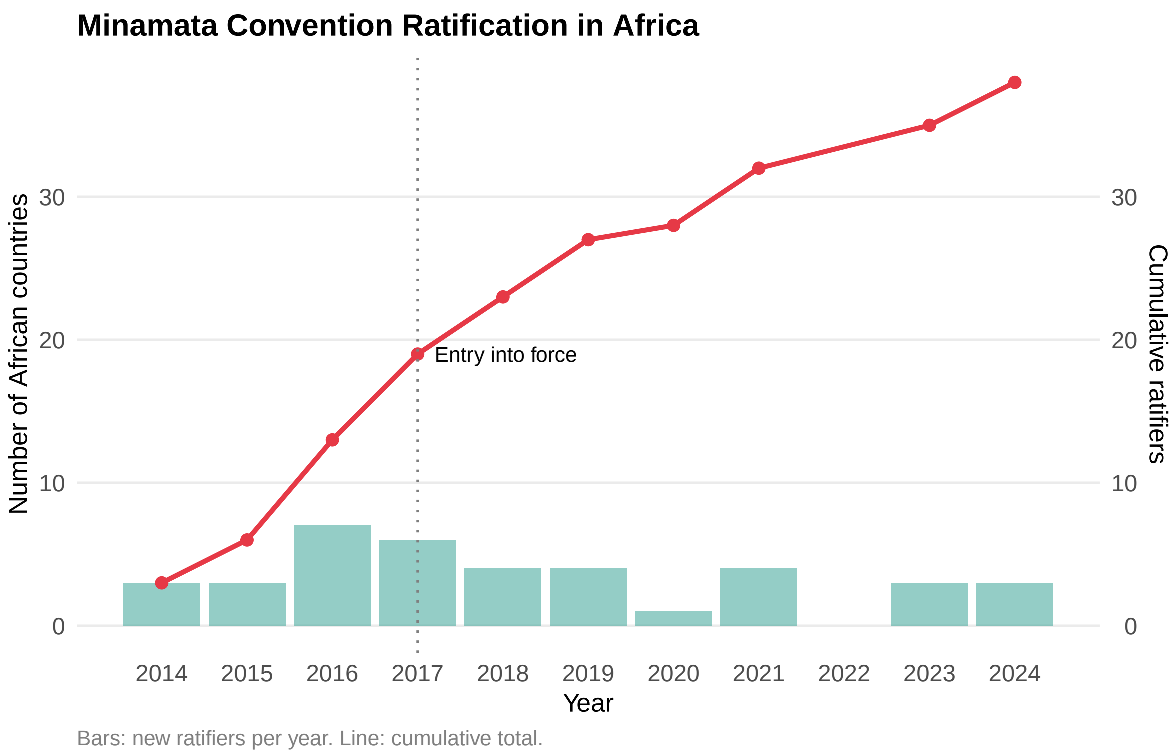

Figure 7: Minamata Convention Ratification in Africa. Bars show new ratifiers per year. Line shows cumulative total. Dotted vertical line: entry into force (2017).

B. Callaway-Sant'Anna ATT Estimates

Table 5: Callaway-Sant'Anna Doubly Robust DiD: Overall ATT

	Never-treated	Not-yet-treated
Overall ATT	-1.6406	0.5296
	(1.6770)	(1.8914)
Estimation method	Doubly Robust (DR)	
Covariates	Log GDP per capita	
Bootstrap iterations	1,000	
Treatment cohorts	7	
Countries	54	

Notes: Callaway and Sant'Anna (2021) group-time ATT estimates aggregated to overall ATT. Doubly robust estimation combines inverse probability weighting with outcome regression. Treatment: first full calendar year after Minamata Convention ratification. Standard errors from multiplier bootstrap (1,000 iterations).

Table 6: Minamata Convention Ratification: African Countries

Country	Year	NAP	ASGM	Country	Year	NAP	ASGM
Djibouti	2014	—	—	Niger	2018	2022	high
Gabon	2014	2024	—	Nigeria	2018	2021	moderate
Lesotho	2014	—	—	Guinea	2018	2021	high
Seychelles	2015	—	—	Guinea-Bissau	2018	—	—
Mauritania	2015	—	—	Uganda	2019	2021	moderate
Chad	2015	2022	—	South Africa	2019	—	—
Senegal	2016	2019	high	Côte d'Ivoire	2019	2023	moderate
Zambia	2016	2023	moderate	Equatorial Guinea	2019	—	—
Mali	2016	2020	very$_high$	Tanzania	2020	2022	very$_high$
Botswana	2016	—	—	Cameroon	2021	2024	moderate
Sierra Leone	2016	2020	high	Burundi	2021	2020	—
Benin	2016	—	low	Central African Republic	2021	2021	moderate
Gambia	2016	—	—	Zimbabwe	2021	2021	high
Togo	2017	2023	low	Eritrea	2023	2023	—
Ghana	2017	2022	very$_high$	Malawi	2023	—	—
Burkina Faso	2017	2020	very$_high$	Kenya	2023	2022	moderate
Rwanda	2017	2023	—	Mozambique	2024	2024	moderate
Namibia	2017	—	—	Ethiopia	2024	—	moderate
Mauritius	2017	—	—	Liberia	2024	—	moderate

Notes: Year = ratification year. NAP = National Action Plan submission year. ASGM = artisanal mining prevalence level.

C. Minamata Convention Ratification Details

D. Data Appendix

D.1 Mercury Trade Classification

Mercury metal is classified under HS code 280540 in the Harmonized Commodity Description and Coding System. This 6-digit code specifically covers mercury (quicksilver), excluding amalgams. We use import-reported bilateral trade data from the Observatory of Economic Complexity (OEC), which curates UN Comtrade records using the BACI reconciliation methodology developed by CEPII.

The BACI methodology reconciles discrepancies between importer- and exporter-reported flows by weighting reports inversely by their estimated reliability. This is particularly important for mercury trade, where discrepancies between mirror flows are systematically large—a feature documented by **?** as evidence of informal or unreported trade.

D.2 Gold Export Classification

Gold exports are captured under HS heading 7108, which includes three 6-digit subcodes: 710811 (gold powder), 710812 (unwrought gold, non-monetary), and 710813 (semi-manufactured gold, non-monetary). Code 710820 (monetary gold) is excluded because it represents central bank transactions unrelated to mining output. We aggregate across the three subcodes to obtain total gold export values.

D.3 Country Sample

Our sample comprises all 54 African Union member states. We include island nations (Comoros, Cape Verde, Mauritius, São Tomé and Príncipe, Seychelles) and North African states (Algeria, Egypt, Libya, Morocco, Tunisia) despite their minimal ASGM activity, as they serve as natural controls. South Sudan is included from its independence year (2011). Somalia and Eritrea have limited trade data availability.

D.4 Balanced Reporter Definition

A country is classified as a "balanced reporter" if it has positive mercury imports in at least 50% of sample years during 2005–2020 (i.e., at least 8 of 16 years). This threshold is arbitrary but balances data quality against sample size. In our data, 31 of 54 countries meet this criterion.

D.5 EU Member States (2010)

The pre-ban EU mercury share is computed using imports from the following 27 EU member states as of 2010: Austria, Belgium, Bulgaria, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, Netherlands, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden, and the United Kingdom (included as a pre-Brexit member).

E. Identification Appendix

E.1 Pre-Trend Tests for EU Ban

The formal pre-trend test regresses log mercury imports on the interaction of EU mercury share with a linear year trend, restricting to the pre-ban period (2005–2010). The coefficient on EUShare $\times$ Year is small and insignificant ($p = 0.67$), confirming no differential trends between high- and low-EU-dependent countries before the ban.

E.2 Callaway-Sant'Anna Technical Details

We use the `did` R package (version 2.1+) to implement the **?** estimator. Key choices:

- **Estimation method:** Doubly robust (DR), combining IPW with outcome regression (**?**)
- **Control group:** Never-treated (primary), not-yet-treated (robustness)
- **Base period:** Universal (all pre-treatment periods averaged)
- **Pre-treatment covariates:** Log GDP per capita
- **Inference:** Multiplier bootstrap with 1,000 iterations
- **Treatment cohorts:** 8 ratification year groups spanning 2015–2022 (countries ratifying in 2023–2024 are coded as untreated because their first treated year exceeds the data endpoint); 7 of 8 produce estimable group-time ATTs

The estimator requires a balanced panel within the analysis window; two country-year observations with missing outcomes were dropped during estimation. Countries ratifying in 2023–2024 (Eritrea, Kenya, Malawi, Ethiopia, Liberia, Mozambique) are coded as untreated because their first treated year exceeds 2023 (the data endpoint). Warnings about small group sizes (cohorts 2015, 2016, 2019, 2020, 2021, 2022) reflect the fact that some ratification years have only 1–3 countries. The singular covariance matrix warning prevents computation of the pre-test Wald statistic; we rely on visual inspection of the event study plot for pre-trend assessment.

E.3 Wild Cluster Bootstrap

With 54 country clusters, asymptotic cluster-robust standard errors may over-reject. We report a Wald test statistic of 6.32 ($p = 0.012$) for the EU ban coefficient, computed with country-clustered variance. This provides reasonable confidence in the significance of the EU ban result despite the moderate number of clusters.

F. Robustness Appendix

F.1 Extensive Margin

The linear probability model for any mercury imports (binary) yields a coefficient of -0.271 on EU Share $\times$ Post, indicating that EU-dependent countries were 27 percentage points less likely to report any mercury imports after the ban. This confirms that the EU ban affected both the intensive margin (import values conditional on importing) and the extensive margin (the probability of importing at all).

F.2 Transit Hub Analysis

Togo and South Africa are known mercury transit points—countries that import mercury primarily for re-export to neighboring countries rather than domestic ASGM use. Their inclusion might bias estimates if the EU ban's primary effect was to redirect transit flows. Excluding both countries yields a coefficient of -2.607, essentially unchanged from the baseline (-2.464), indicating that the result is not driven by transit dynamics.

F.3 Alternative Transformation

The inverse hyperbolic sine (IHS) transformation $\text{IHS}(x) = \ln(x + \sqrt{x^2 + 1})$ approximates the log for large values while preserving zeros without the arbitrary $\ln(x + 1)$ addition. The IHS coefficient of -2.650 is slightly larger than the log+1 baseline, suggesting that the log+1 transformation, if anything, attenuates the estimate.