Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Measuring the Effect of Tariffs on Critical Minerals Trade in a Multidimensional Panel

Muhammad Afnan Arif

muhdafnanarif@gmail.com

Asia-Europe Institute, Universiti Malaya

Research Article

Keywords: multi-dimensional panel data, structural gravity, non-identification, Poisson pseudo-maximum likelihood, critical minerals

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11503075/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Measuring the Effect of Tariffs on Critical Minerals Trade in a Multidimensional Panel

Muhammad Afnan Arif*
*Asia-Europe Institute, Universiti Malaya, Kuala Lumpur, Malaysia.
Corresponding author: muhdafnanarif@gmail.com.

Abstract

Estimated tariff elasticities for critical minerals change several-fold as fixed effects are added and this instability is usually treated as specification sensitivity. This paper traces it to non-identification. A tariff schedule published for an importer, a product and a year takes one value for every supplier, so in a structural gravity system with importer-product-year effects the rate cancels from the gravity equation and its coefficient describes a parameter absent from the model. The remedy is a rate that differs across suppliers through preferential access. A Poisson pseudo-maximum likelihood gravity model with exporter-mineral-year, importer-mineral-year and pair-mineral effects on a multi-dimensional panel of bilateral mineral trade gives a negative tariff elasticity, significant under pair-clustered inference, much smaller than additive-effects estimates and little changed by a bias correction. Log-linear, nonparametric and interactive-effects estimators find no significant pooled response. Trade in cobalt, copper and graphite falls with the tariff, trade in lithium rises and the rare earths show no response. The residual variance share shows that few flows carry a supplier-specific rate, so the estimate is an association identified from few cells and not a general elasticity. It is not significant within short subperiods and the interval includes zero under a pair bootstrap.

Keywords: multi-dimensional panel data; structural gravity; non-identification; Poisson pseudo-maximum likelihood; critical minerals

JEL Classification: C23, F13, F14

1. Introduction

Demand for the minerals on which clean energy and digital technologies depend is projected to grow quickly while their supply remains concentrated. The United Nations Conference on Trade and Development (UNCTAD, 2026b) calculates from International Energy Agency data that projected demand rises by 353% for lithium and by 131% for graphite between 2024 and 2040 and reports that in 2025 the Democratic Republic of the Congo mined 74% of the world's cobalt and China 78% of its natural graphite. Refining is similarly concentrated, with the average share of the largest refining country in the key minerals other than rare earths rising from 70% in 2023 to 72% in 2025 (International Energy Agency, 2026). Critical minerals have accordingly become a policy priority (Ali et al., 2017; Gulley et al., 2018; Kalantzakos, 2020; Vivoda, 2023; Watari et al., 2020).

Governments hold two levers over trade in these minerals. On the export side, UNCTAD (2026b) counts nearly 100 measures on critical energy transition minerals introduced since 2020, among them 37 licensing requirements, 31 export taxes and 29 export bans. On the import side, the most-favoured-nation (MFN) tariff applied to energy-related critical minerals averaged about 4% in 2022 while bound rates average about 26% (Snoussi-Mimouni & Avérous, 2024), which leaves importers wide room to raise duties. The United States applies a 50% tariff to semi-finished copper products, on their copper content from 1 August 2025 and on their full value from 6 April 2026 while copper ores, cathodes and anodes are exempt (Congressional Research Service, 2026).

Preferential access is the part of import tariff policy that differs across suppliers to the same market and it is spreading: UNCTAD (2026a) identifies 73 international agreements and partnership instruments on critical minerals, 58 of them signed after 2022, although not every one carries tariff preferences. In the panel studied here, trade in cobalt, copper, graphite,

lithium and the rare earths among the 50 largest exporters and the 50 largest importers grew from USD 20.6 billion in 1995 to USD 199.1 billion in 2024 in current prices, with copper accounting for 92.4% of the value over the whole period. Figure 1 places these minerals among other minerals traded worldwide, where they rise from 31% of export value in 1995 to 50% in 2024 and shows the exporters inside and outside the panel, the latter supplying 13% of the value in 2024. Whether tariffs shape this trade, and by how much, is therefore a question with direct policy weight.

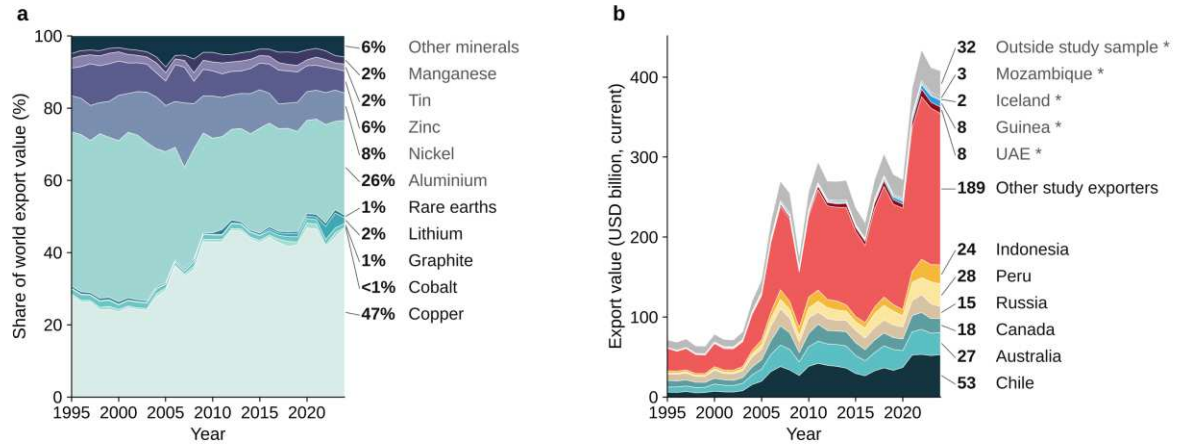


Figure 1. World exports by mineral and by exporting country.

Notes: Panel (a) plots the share of world export value by mineral and panel (b) the export value in USD billion at current prices by exporting country, both for 1995 to 2024 from the Base pour l'Analyse du Commerce International (BACI) database in the 1992 Harmonized System (HS92), version 202601, the five minerals of the study are copper, cobalt, graphite, lithium and the rare earths, labelled in bold, and cover the six-digit lines of Table D5 that appear in the trade data, the other minerals are aluminium, nickel, zinc, tin, manganese and a group of titanium, chromium, tungsten, molybdenum, niobium, tantalum, vanadium, zirconium and antimony ores, defined for this figure and listed by six-digit line in the replication package, the labels give the share in % in 2024 in panel (a) and the value in USD billion in 2024 in panel (b) with less than 1% for a share below 0.5%, exporters marked * are not among the 50 exporters of the panel, other study exporters are the remaining exporters of the panel and outside study sample is every other exporter including BACI aggregates.

A small empirical literature has assessed how export restrictions affect trade in these minerals (Korinek & Kim, 2011), but tariff elasticities estimated for the panel studied here are unstable across specifications. A coefficient that changes by a factor of several as fixed effects (FE) are added is usually read as ordinary specification sensitivity and the specification with the most FE is preferred. This paper argues that the instability has a different source and a different remedy.

The source lies in how the regressor varies. Product-level tariffs are conventionally recorded at the importer, product and year margin, the level at which an MFN schedule is published, so the rate takes one value for every origin serving a market. In a structural gravity system a trade cost that is invariant across origins factors out of the inward multilateral resistance (MR) term and cancels from the gravity equation. It has no bilateral parameter and its residual variance share (RVS), the proportion of its variance that remains after the FE required by the model are removed, is exactly zero. Any elasticity reported for it reflects numerical properties of the estimation algorithm and not a structural parameter, which explains the instability.

The remedy is to use rates that vary across origins. Preferential access does not cover every exporter to a market in a given year, so an exporter with access can face a rate below the MFN rate applied to other partners. This paper constructs a realised bilateral tariff (RBT), built from the partner-specific rate (PSR) where one is recorded and the MFN rate otherwise and entered as the log of one plus the rate in %. The RBT retains a positive RVS where the MFN rate retains none, but the variation is limited because only a small share of observations

carries a PSR. Boehm et al. (2023) use the same origin logic, exploiting changes in MFN tariffs that reach only the partners trading on those terms to estimate trade elasticities at short and long horizons. Fontagné et al. (2022, p. 8) note that in their product-by-product regressions the MFN tariff is captured by the importer-time effect, so that the tariff coefficient is identified from variation across importers for each exporter and product.

With the remedy in place, the elasticity is estimated by PPML-HDFE, that is, Poisson pseudo-maximum likelihood (PPML) with high-dimensional fixed effects (HDFE; Correia et al., 2020; Santos Silva & Tenreyro, 2006), with exporter-mineral-year (EMT), importer-mineral-year (IMT) and pair-mineral (PM) effects, the structure that the theory requires. It is negative and statistically significant under pair-clustered standard errors (SEs), much smaller in absolute magnitude than in the additive-effects specifications of the conventional sequence and of similar size to specifications that include a mineral dimension. By mineral, the elasticity is negative for cobalt, copper and graphite, positive for lithium and indistinguishable from zero for the rare earths. Because most of the information on the rate comes from copper, the pooled estimate lies close to the copper elasticity. Log-linear estimators, the nonparametric estimator of Henderson et al. (2022) and the interactive fixed-effects estimators of Kapetanios et al. (2021) and Jin et al. (2025), all with the same FE, find no significant pooled response on the same data or on a balanced block of it, whereas an analytic bias correction (Weidner & Zylkin, 2021) leaves the Poisson estimate nearly unchanged.

Turning to scope, tariffs are one of several measures that bear on trade in these minerals and the paper makes no claim about their share of the policy mix. Export restrictions, the main alternative, act on the exporter side, but the restriction index in the data is non-zero in too few cells to support an estimate and a restriction applied to every destination is, like the MFN tariff, constant within a cell that the EMT effects absorb. Preferential tariffs are the only measure in the data that differs across suppliers to the same market, which is what identification in this setting requires.

The paper asks whether a tariff that is invariant across origins can be identified in a structural gravity model with importer-mineral-year effects, what elasticity results when identification comes from the RBT and how reliable it is. It reports the RVS beside every coefficient, estimates the elasticity pooled and by mineral, sets it against alternative estimators and assesses its reliability through simulation, subperiod, lead, coverage-restricted and bootstrap checks.

The paper makes three contributions. That variables affecting imports without regard to origin, such as a country-level applied tariff, cannot be identified once importer effects are included is a familiar point in the gravity literature (Head & Mayer, 2014). The first contribution is to trace that point through the structural gravity system, where a tariff collected at the importer and product level cancels from the gravity equation itself and the instability of the estimate across specifications follows. The second is to report the RVS beside every coefficient so that a reader can see how much variation an estimate rests on. The RVS is elementary in construction, being one minus the coefficient of determination of an auxiliary regression and the contribution is the reporting practice built around it and the boundary case in which it is exactly zero. The third is to apply both to critical minerals, to set the Poisson estimate beside established estimators that run on the same data, to examine effect structures shared across minerals and to state how far the result can be trusted.

The reliability results qualify the third contribution. The estimate is not significant within short subperiods, in samples restricted to importers and years with partner records or under a pair bootstrap. Future rates also predict current trade. It is therefore read as a narrow estimate of an association and not as a causal or general elasticity.

Section 2 reviews the related literature. Section 3 sets out the model, derives the cancellation result, defines the RVS, describes the estimation strategy and presents the data and the construction of the RBT. Section 4 reports the results and examines their reliability and Section 5 concludes. Derivations, the data construction and further results are in the appendices. An

online supplement provides the proofs of the results stated in the appendices, a Monte Carlo study and additional empirical results.

2. Related literature

The foundations of structural gravity run from the origin-differentiated demand of Armington (1969) and Bergstrand (1985) through the Ricardian formulation of Eaton and Kortum (2002) and the firm-selection accounts of Chaney (2008) and Helpman et al. (2008). Empirical methods are set out by Anderson (1979), Feenstra (2004) and Yotov et al. (2016); Head and Ries (2001) measure border effects in US-Canada trade. PPML is the standard estimator because trade contains many zeros and the estimator remains consistent under the heteroskedasticity typical of trade data (Santos Silva & Tenreyro, 2006, 2011) and estimation with high-dimensional FE is now routine (Correia et al., 2020; Larch et al., 2019). Exporter-time and importer-time effects account for MR (Baldwin & Taglioni, 2007) and the fixed effects of PPML satisfy the adding-up constraints that define the structural MR terms when no observations are missing (Fally, 2015) while pair effects absorb time-invariant bilateral frictions, including those correlated with agreement membership (Baier & Bergstrand, 2007). Two limitations are documented. Weidner and Zylkin (2021) characterise the incidental-parameter bias of PPML with three-way FE and their analytic correction is applied below. Pfaffermayr (2026) treats the downward bias of pair-clustered SEs, which the reliability analysis addresses through alternative clustering and a pair bootstrap.

Closest to the application, Fontagné et al. (2022) estimate product-level tariff elasticities in a gravity framework, with the tariff entering as the logarithm of one plus the ad valorem rate and Boehm et al. (2023) identify trade elasticities at short and long horizons from changes in MFN tariffs that reach only the partners trading on those terms. The present paper differs in tracing the consequence of origin invariance through the structural system and in reporting what identifying variation each specification leaves.

Turning from the tariff literature to the fixed-effect structure, Egger and Pfaffermayr (2003) added the time-invariant exporter-importer interaction effect to the exporter, importer and time main effects of the three-way gravity panel. Baltagi et al. (2003) added the exporter-time and importer-time interaction effects, which gives the full set of three interaction effects that the structure used here generalises by adding the mineral dimension. Balázsi et al. (2018) derived the corresponding within estimators, including for unbalanced and no-self-flow data. Their models are linear with additive effects and their four-dimensional model contains the structure used here as the special case without a pair-year effect. They show that their no-self-flow transformations do not work when the exporter and importer sets overlap only in part, as in the present sample. They note that no obvious transformation can be worked out for that case and base their treatment of unbalanced data instead on projections computed on the observations present, which is the route followed here. They do not treat Poisson estimation, zeros or the question of which regressors remain after the projection. Davis (2002) obtained the same projection algebra for error components models with any number of effects on unbalanced data, again for linear models with additive effects.

A further strand treats multi-dimensional panels with richer error structures. Kapetanios et al. (2021) use cross-section averages as proxies for global and local common factors. Jin et al. (2025) estimate the global and local factors themselves and let the slopes differ by pair. Lu and Su (2026) show that for balanced three-dimensional panels with additive effects the within, Mundlak and least squares dummy variable estimators coincide. These estimators are developed for linear regression models, so they are run below on a balanced block with the log of one plus exports.

Beyond linear and Poisson specifications, Henderson et al. (2022) propose a closed-form nonparametric estimator of the gradient of an unknown conditional mean for panels with additive fixed effects of any dimension. It double-demeans the data, differences over all pairs of periods and weights each pair by a kernel in the regressor with a cross-validated bandwidth.

Its three-dimensional case with pair, exporter-time and importer-time effects is the structure used here without the mineral dimension. Because its transformations are defined on a complete array and its fixed effects enter additively, so that it does not nest the multiplicative gravity equation or its zero flows, the estimator is run below on a balanced block as a complementary check and not as the main estimator.

Moment-based alternatives to FE dummies exist for multiplicative models. Jochmans (2017) derives moment conditions that are free of fixed effects for a two-way model on a single cross-section of directed dyads and Yang and Zhang (2023) extend that strategy to three-way gravity models with multiplicative unobserved effects. That route is not followed here for three reasons. For a regressor that varies only by importer, as an origin-invariant tariff does in a single cross-section, the population moment of Jochmans (2017) is zero at every parameter value whatever the instrument function, which mirrors the cancellation by the IMT effect and offers no route around it. The realised rate behind the RBT differs from the MFN rate only on PSR rows, so a moment of this form would take its information from those rows alone. The author is not aware of a version for a common slope across the mineral-specific three-way panels that the four-dimensional structure comprises. Jochmans (2017) also reports that direct evaluation grows with the fourth power of the number of agents and that the Poisson variance estimator tends to deliver SEs that are too small, which is one reason the inference below is pair-clustered and checked by a pair bootstrap.

A separate literature chooses the FE structure by data-driven selection. Post-selection inference (Belloni et al., 2014a, 2014b), resting on the penalised regression of Tibshirani (1996), supplies valid confidence intervals after a data-driven choice among many controls and Chiang et al. (2023) extended it to three-dimensional panels in which the candidate controls are themselves sets of FE. A selection procedure searches over a candidate set chosen before any data are examined. If the structure a regressor requires lies outside it, the RVS is zero and the procedure never evaluates the effect that would absorb the regressor, so inference is conducted for a parameter that the estimating equation does not identify. Chiang et al. (2023) note in a footnote that the coefficients of certain controls, such as exporter-year and importer-year gross domestic product when exporter-year and importer-year effects are included, are not identified under more general combinations of FE whereas the coefficients of dyad-level regressors such as distance and trade agreements are identified under all the specifications they consider. This paper takes the non-identified case as its subject. Sasaki and Ura (2026) treat a related problem, inference on average partial effects of slow-moving regressors in correlated random coefficient models for short panels.

3. Model and data

3.1 Model

Gravity is used here to settle which effects a bilateral tariff elasticity requires, because a multi-dimensional panel describes how the data are indexed and not which model applies to them. Structural gravity delivers an estimating equation in which bilateral trade depends on the bilateral trade cost and on two MR terms that summarise the trading opportunities available to each side (Anderson & van Wincoop, 2003). When trade is observed by product, the value of trade from exporter i to importer j in mineral m and year t satisfies:

$$X_{ijmt} = \frac{Y_{imt} E_{jmt}}{Y_{mt}} \left(\frac{\tau_{ijmt}}{\Pi_{imt} P_{jmt}} \right)^{1-\sigma_m}$$

where Y_{imt} is the shipments of the exporter, E_{jmt} the expenditure of the importer, Y_{mt} world output, τ_{ijmt} the bilateral trade cost, Π_{imt} and P_{jmt} the outward and inward MR terms and σ_m the elasticity of substitution. The trade cost combines the applied tariff θ_{ijmt} with freight, insurance, time and every other friction, $\tau_{ijmt} = \left(1 + \theta_{ijmt}\right) \kappa_{ijmt}$. The full system with the resistance terms and market clearing is in Appendix A.

Taking logarithms, the terms that depend on one side of the market only are collected into EMT and IMT effects, which account for MR (Baldwin & Taglioni, 2007) and are recovered exactly by PPML (Fally, 2015) and a directed PM effect absorbs time-invariant bilateral frictions, including those correlated with agreement membership (Baier & Bergstrand, 2007). The estimating equation is:

$$\mathbb{E}\left[X_{ijmt} \mid \alpha_{imt}, \gamma_{jmt}, \eta_{ijm}, x_{ijmt}\right] = \exp\left(\alpha_{imt} + \gamma_{jmt} + \eta_{ijm} + x_{ijmt}'\beta\right)$$

where x_{ijmt} collects the bilateral policy variables with parameters β, α_{imt} and γ_{jmt} are the EMT and IMT effects and η_{ijm} is the PM effect. In the application x collects the RBT, constructed below, AGR, an indicator that the pair shares a trade agreement, and four gravity covariates, namely DIST, the log of bilateral distance in km, and the indicators CONTIG for a shared border, LANG for a shared official language and COL for a shared coloniser.

Which tariff measures can carry a bilateral parameter depends on how the measure varies across origins. An MFN tariff takes the same value for every origin and for such a variable the identification failure is not a matter of degree. Write the trade cost as $\tau_{ijmt} = T_{jmt}\,\kappa_{ijmt}$, where T_{jmt} varies only with the importer, the mineral and the year. The inward MR term is a weighted sum of bilateral trade costs over all origins (Freeman et al., 2025), so the origin-invariant component can be taken outside the summation:

$$P_{jmt}^{1-\sigma_m} = T_{jmt}^{1-\sigma_m} A_{jmt}, \quad A_{jmt} = \sum_i \left(\frac{\kappa_{ijmt}}{\Pi_{imt}}\right)^{1-\sigma_m} \frac{Y_{imt}}{Y_{mt}}$$

Substituting the expression for the inward resistance term into the gravity equation gives:

$$X_{ijmt} = \frac{Y_{imt}E_{jmt}}{Y_{mt}} \left(\frac{T_{jmt}\kappa_{ijmt}}{\Pi_{imt}P_{jmt}}\right)^{1-\sigma_m} = \frac{Y_{imt}E_{jmt}}{Y_{mt}} \frac{\kappa_{ijmt}^{1-\sigma_m}}{\Pi_{imt}^{1-\sigma_m} A_{jmt}}$$

The factor T_{jmt} appears in the numerator through the bilateral trade cost and in the denominator through the inward MR term, in the same power, so it cancels and drops from the outward MR term as well (see Appendix B). The conclusion is stronger than absorption, because a trade cost common to all origins raises the delivered price of every source in the same proportion and relative demand is unchanged. What remains is a general-equilibrium effect through the aggregates that respond to the tariff, the shipments of the exporter and the expenditure of the importer, which are indexed on exporter, mineral and year or on importer, mineral and year and are absorbed by the EMT and IMT effects by construction.

By contrast, Freeman et al. (2025) recover a country-time effect from an identity that links it to output and multilateral resistance, but no such identity exists for an IMT effect with a tariff inside it, because the tariff cancels from the system and cannot be recovered from it. Two conditions are required and both hold in the application. The measure must be constant across every origin in the sample, which is verified in the data rather than assumed and the sample must contain no domestic flows, since a tariff that falls on imports but not on domestic sales discriminates and does not cancel. Where either fails the cancellation is partial and the RVS in Table D4 reports how partial.

Cancellation is the extreme case of a wider property of the fixed effects, since the parameters are identified only from variation in x that remains after all three sets of effects are removed. In four dimensions the requirement is severe, because an effect indexed on exporter, mineral and year partitions the sample far more finely than one indexed on exporter and year and a variable that moves freely across exporters and years can still be constant within every importer, mineral and year cell.

Let the index set contain the exporter, the importer, the mineral and the year. A fixed-effect structure F is a collection of subsets of that index set, each defining one set of effects, so that a subset containing exporter, mineral and year defines an effect taking a distinct value in every exporter-mineral-year cell. In the structure used here every subset contains the mineral, so no

effect is shared across minerals, the five minerals are separate gravity panels that share the slope and the residual variation of the regressor is computed within each mineral.

Removing these effects from the data is a projection. With D_F the matrix that stacks the indicator columns of every subset, the residual maker $\mathbf{M}_F = I - D_F(D_F{}'D_F)^{-}D_F{}'$ projects onto the orthogonal complement of the design. For a complete array it has the closed form $\mathbf{M}_F = Q_{N_i} \otimes Q_{N_j} \otimes I_{N_m} \otimes Q_{N_t}$, with $Q_N = I_N - N^{-1}\iota_N\iota_N{}'$, where I_N is the identity matrix and ι_N a vector of ones of order N (see Appendix A). The estimation sample is not a complete array, because own-country flows are absent and separated cells are dropped. Balázsi et al. (2018) show that within transformations derived for complete arrays are biased on such data whereas a projection computed on the rows actually used is not, which is what alternating-projection routines deliver (Correia et al., 2020). Because structures of interest are nested, the variation available to identify any parameter is non-increasing as the structure tightens.

The variation profile of a regressor, the indices along which it varies, is a property of the realised data and not of the name the variable carries. A tariff described as bilateral may be applied on an MFN basis and take one value for all origins and a concentration measure described as exporter-specific may be constructed within exporter-mineral-year groups. A regressor is measurable with respect to a subset f of the indices when it can be written as a function of the indices in f alone. If the structure contains an effect indexed on such a subset, the indicator columns of that effect span the regressor and residualising leaves nothing behind. Measurability is sufficient but not necessary, because the RVS is zero if and only if the regressor lies in the column span of the design, which also covers sums of such functions, for example one term indexed on exporter, mineral and year plus one indexed on importer, mineral and year.

The RVS of a regressor under structure F is the ratio of the variance of the residualised regressor to the variance of the regressor in deviation from its overall mean,

$$\rho(x, F) = \frac{\mathrm{Var}(\mathbf{M}_F\tilde{x})}{\mathrm{Var}(\tilde{x})} = \frac{\tilde{x}\prime\mathbf{M}_F\tilde{x}}{\tilde{x}\prime\tilde{x}} = 1 - R^2(\tilde{x} \text{ on } D_F)$$

where $\tilde{x}$ is the centred regressor. The RVS equals one minus the coefficient of determination of an auxiliary regression of $\tilde{x}$ on the design, is bounded between zero and one and is non-increasing as the structure tightens. It is computed by alternating projections without forming the dummy matrix (Correia et al., 2020) and on the unweighted design rather than the design weighted by the Poisson objective, because absorption is a property of the relation between a regressor and a partition of the sample and not of a likelihood evaluated at particular parameter values (see Appendix C).

The RVS is not a test of whether inference is valid and it bounds identifying variation, not bias: a large share is compatible with endogeneity and a small share with exogeneity. An RVS of zero means the parameter is not identified under the structure imposed, so the question is not whether a critical value is reliable but whether there is a parameter to attach one to. Routines for HDFE estimation detect exact collinearity with the effects (Correia et al., 2020) but not near collinearity, so a regressor that retains a small RVS is reported with a standard error that does not reflect how little variation identifies it. Intermediate values are the informative case, because the coefficient is then identified from a stated fraction of the variation and should be read against that fraction.

For estimation, the estimating equation is fitted by PPML-HDFE, computed by alternating projections (Correia et al., 2020). PPML is used because trade contains many zeros and the estimator remains consistent for the conditional mean under the heteroskedasticity typical of trade data (Santos Silva & Tenreyro, 2006, 2011). Zeros are retained and observations that the FE absorb as singletons or separate are dropped. The score equations, the weighted within transformation and the variance estimator are in Appendix A. The alternatives, applied to the same regressor and sample, are ordinary least squares (OLS) of the log of one plus exports on all rows (OLS-L1P) and of the log of positive exports (OLS-LN), Gamma pseudo-

maximum likelihood on positive flows (GPML), the PPML estimate corrected for incidental-parameter bias (PPML-BC) by the analytic formula of Weidner and Zylkin (2021), the nonparametric estimator of Henderson et al. (2022), the three-dimensional common correlated effects (3DCCE) estimator of Kapetanios et al. (2021) and the multi-level interactive fixed-effects (IFE) estimator of Jin et al. (2025). Both OLS variants apply the within estimator of Balázsi et al. (2018) to a transformation of exports and are generally inconsistent for the elasticity under heteroskedasticity, because the expectation of the logarithm of the error depends on the regressors (Santos Silva & Tenreyro, 2006).

The multi-dimensional model of the paper is PPML-HDFE with EMT, IMT and PM effects, a Poisson model of exports on the four indices exporter, importer, mineral and year with the RBT as the regressor of interest. The nonparametric, 3DCCE and IFE estimators are defined for three indices, so they are applied within each mineral on the balanced block while the remaining estimators are applied to the same four-index array under the same effects.

The extension relaxes the restriction that no effect spans two minerals. The pooled structures drop the mineral from the exporter-year, importer-year or pair effects in turn). A pooled effect restricts the mineral-specific effect to take the same value for every mineral, so it is valid only if the multilateral resistance of an exporter or an importer and the time-constant bilateral relation do not differ across minerals. The RVS of the regressor is computed under each structure and shows how much variation identifies the elasticity once minerals are pooled.

SEs are clustered on the directed pair, which leaves serial dependence within a pair unrestricted. They are not corrected for the downward bias of pair-clustered SEs treated by Pfaffermayr (2026). Every effect of the EMT, IMT and PM structure is specific to a mineral, so column (8) of Table 1 stacks five three-way gravity panels that share one slope. Within a panel the incidental-parameter bias is characterised by Weidner and Zylkin (2021), so their correction is applied to each mineral and the five corrections are stacked for the common slope, an adaptation by the author that the original paper does not provide. The SEs of the corrected estimates are those of the uncorrected fit.

The nonparametric estimator is defined on a complete array with additive fixed effects, so it is applied to the log of one plus exports on the largest complete block of the estimation sample, with pair-mineral, exporter-mineral-year and importer-mineral-year effects, a cross-validated bandwidth, AGR entering linearly and SEs from a bootstrap over directed pairs. Its average derivative is a gradient of the conditional mean of the log of exports and is set beside the OLS-L1P and PPML estimates on the same block. The implementation is the author’s, written from the formulas of the original paper.

The two interactive-effects estimators are developed for linear regression models, so they are applied to the same block and the same outcome, the log of one plus exports, in which the zero flows enter as zeros. The pooled 3DCCE estimator augments the regression of each pair with cross-section averages of the outcome and the regressors taken over all pairs of the mineral and over the partners of the exporter and of the importer, which proxy for global and local common factors without estimating them. The pair regressions are then combined into one common slope. The multi-level IFE estimator extracts global factors and exporter-specific and importer-specific local factors by principal components and iterates between the factors and a common slope, with the numbers of factors chosen by the growth-ratio criterion. Both are fitted within each mineral and stacked for the pooled column, with SEs that allow for dependence within each exporter and within each importer. For each mineral block the OLS-L1P estimates of Panel D of Table 3 are also the Mundlak estimates, because the within, Mundlak and least squares dummy variable estimators coincide (Lu & Su, 2026), which a direct check on the block confirms to numerical precision. Both implementations are the author’s and are described, with small simulations,).

Three procedures are not reported. The selection estimator of Helpman et al. (2008) needs an exclusion restriction for entry into a market that the data do not contain. The jackknife SEs of Pfaffermayr (2026) were attempted but the refits did not complete. The moment estimator

of Yang and Zhang (2023) assumes that the sample contains every pair, whereas the tariff variables are observed for fewer than half of the pair-year rows because only 26 of the 50 importers in the panel have an applied-tariff record. The pair-clustered intervals should be read with caution because the mean SE falls short of the sampling standard deviation in the Monte Carlo simulation.

3.2 Data

The framework is applied to a panel of bilateral mineral trade built from three open data sources. Trade values come from the Base pour l'Analyse du Commerce International (BACI) database of the Centre d'Études Prospectives et d'Informations Internationales (CEPII; Gaulier & Zignago, 2010). Six-digit products of the Harmonized System (HS) are assigned to five minerals, namely cobalt, copper, graphite, lithium and the rare earths, across extraction, processing and downstream stages (see Appendix D). Gravity covariates come from the CEPII gravity database (Conte et al., 2022) with population-weighted distances (Mayer & Zignago, 2011). Applied tariffs at the six-digit level, both partner-specific and for all partners together, come from the UNCTAD Trade Analysis Information System through the World Integrated Trade Solution (WITS) interface.

The panel is a complete array by construction. Ordered pairs formed from the 50 largest exporters and the 50 largest importers by trade value, excluding self-pairs, are crossed with the five minerals and the 30 years from 1995 to 2024, giving 369,600 cells, most of which record no trade. The zeros are retained as equilibrium outcomes and PPML-HDFE is chosen on the multiplicative structure of the gravity model and the share of zero flows and not on a preliminary test (Santos Silva & Tenreyro, 2022). Table D1 traces the construction of the samples. Sample A has 186,383 cells on which the MFN rate of the gravity panel (MFN-G), AGR and DIST are observed and the bilateral tariff analysis uses the smaller sample E of 159,157 cells on which the RBT, AGR, DIST and MFN-G are all observed. Tables D2 and D3 define and summarise the variables of sample E, Table D5 lists the product lines of each mineral, Table D6 the importers and exporters and Table D7 the rows with a PSR in each year.

Figures 2 and 3 describe the panel before any tariff enters. Figure 2 plots the total exports of the 30 largest exporters of each mineral by year: exports span several orders of magnitude across exporters within each mineral, rise over the period in copper, graphite, lithium and the rare earths and are roughly flat in cobalt. Figure 3 plots the share of cells with zero trade by decile of bilateral distance and year: pooled over the five minerals it falls from 83.8% of cells in 1995 to 73.5% in 2024 and in 2024 it is 44.5% in the closest decile and between 69.0% and 81.5% in the other nine. Zeros are therefore the typical outcome in this panel at all but the shortest distances.

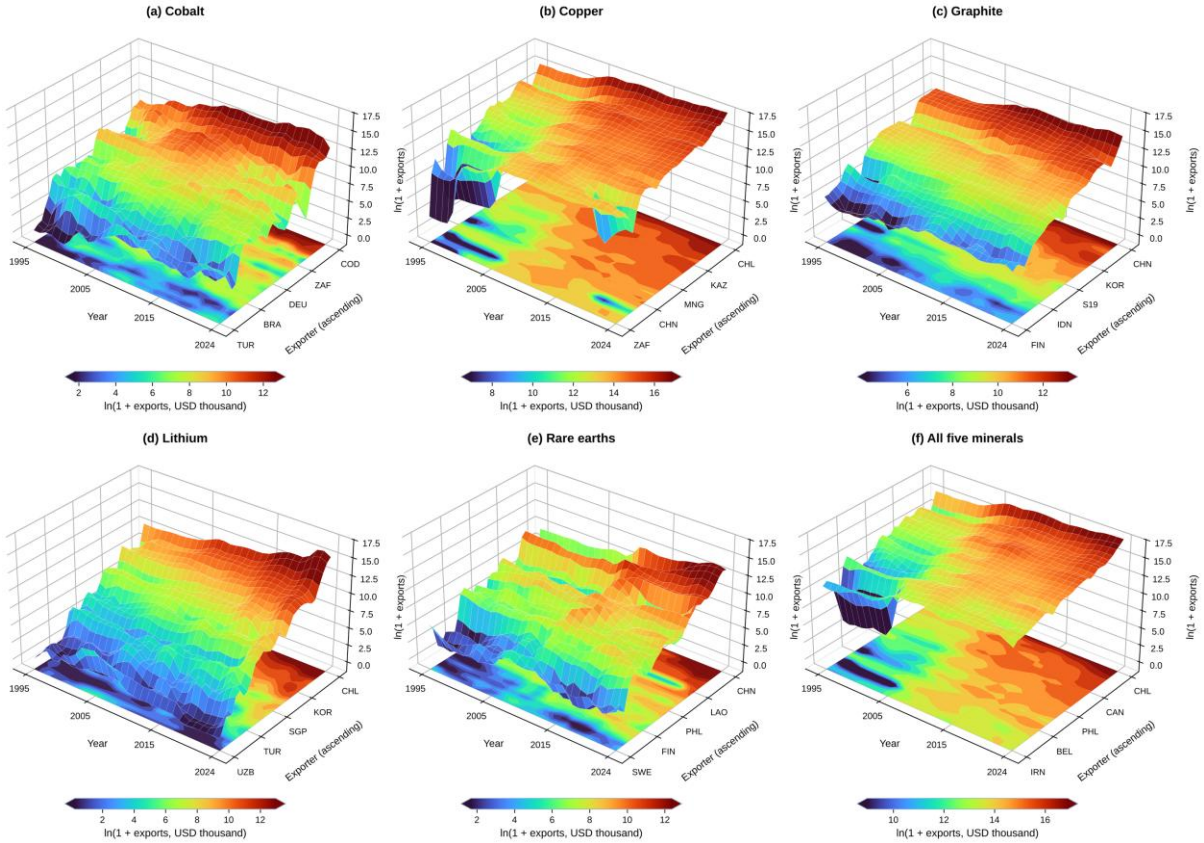


Figure 2. Exports by exporter and year.

Notes: Panels (a) to (e) cover one mineral each and panel (f) pools the five minerals, each panel plots for the 30 exporters with the largest cumulative exports of the mineral in the panel and for each year from 1995 to 2024 the log of one plus total exports to all importers of the panel in USD thousand, exporters are ordered along the front axis from the smallest to the largest cumulative exports and labelled by ISO 3166-1 alpha-3 code except S19, the BACI code for Other Asia, not elsewhere specified, cells with zero total exports are plotted at zero, surfaces are smoothed with a Gaussian kernel of standard deviation 0.7 grid cells, the vertical axis is common to all six panels and the colour bar of each panel spans the 2nd to the 98th percentile of its plotted values.

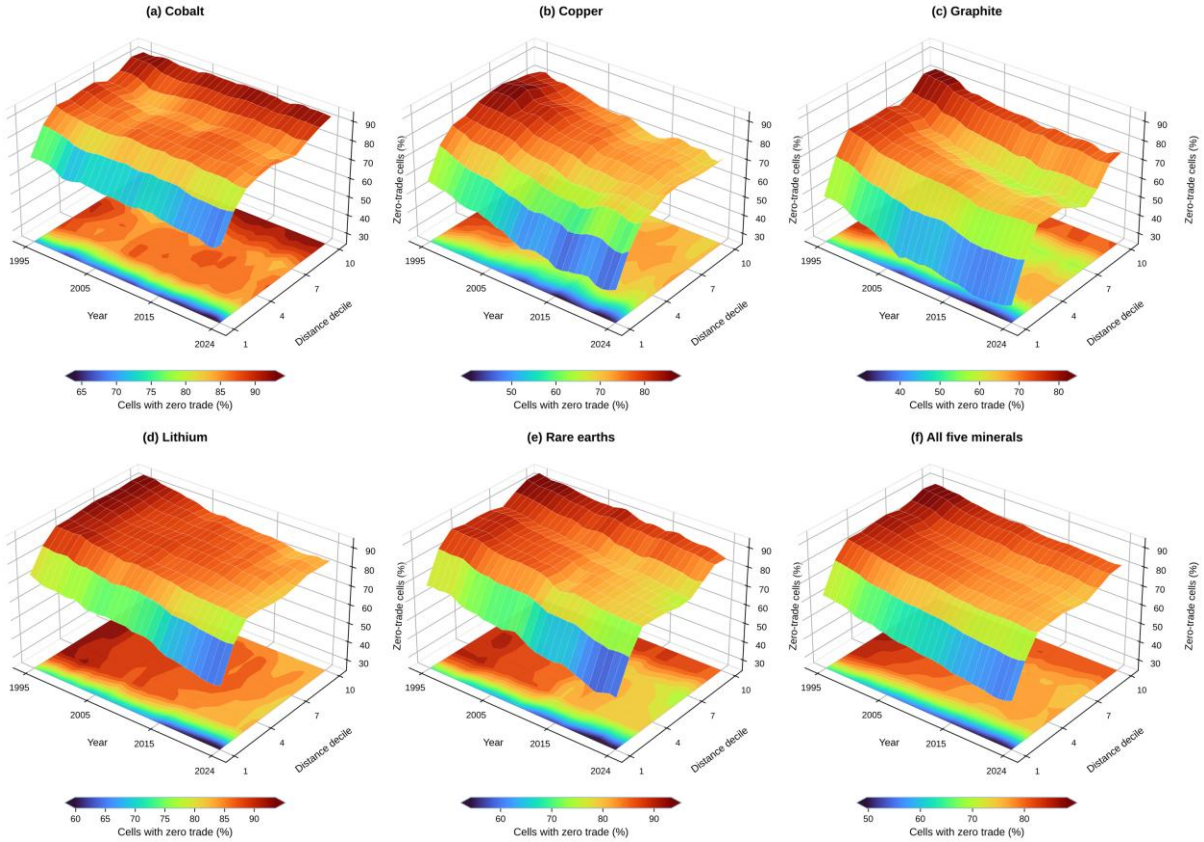


Figure 3. Zero-trade cells by distance decile and year.

Notes: Panels (a) to (e) cover one mineral each and panel (f) pools the five minerals, each panel plots the share of exporter-importer-mineral-year cells with zero trade in each decile of bilateral distance and year in %, deciles are those of bilateral distance in km over the cells of the complete panel with distance observed, decile 1 covers 170 to 1,734 km and decile 10 covers 13,685 to 19,819 km, the 14,700 cells of pairs that include S19 have no distance and are excluded, each decile-year holds 235 to 238 cells in panels (a) to (e) and 1,175 to 1,190 cells in panel (f), surfaces are smoothed with a Gaussian kernel of standard deviation 0.7 grid cells, the vertical axis is common to all six panels and the colour bar of each panel spans the 2nd to the 98th percentile of its plotted values.

For the regressor, the applied rate recorded at the importer, mineral and year level takes a single value for every origin, which is the MFN rate. By the cancellation result its RVS under the required structure is exactly zero, so an elasticity reported for it is a coefficient for a parameter that the model does not contain. Realised bilateral rates depend on whether the importer extends preferential access to the exporter and 94.6% of importer-mineral-year cells in the 186,383-row sample contain both agreement members and non-members. PSRs were retrieved through WITS by six-digit product and year and replace the MFN rate where they exist. Partner groups that are not countries are set aside and no PSR could be retrieved for 2024 or for several importer-year groups, so the RBT is a lower bound on the origin variation in realised rates. The regressor is constructed as:

$$\mathrm{RBT}_{ijmt} = \ln\left(1 + r^{real}_{ijmt}\right), \quad r^{real}_{ijmt} = \mathbf{1}_{ijmt}\, r^{pref}_{ijmt} + \left(1 - \mathbf{1}_{ijmt}\right) r^{mfn}_{jmt}$$

where $\mathbf{1}_{ijmt}$ indicates that a PSR was retrieved, r^{pref}_{ijmt} is that rate in % and r^{mfn}_{jmt} is the MFN rate in %, so that the realised rate is the applied tariff of the structural system measured in % and not as a fraction. Because BACI values are free on board (FOB), the coefficient on the log of one plus the rate as a fraction equals $-\sigma_m$ under constant elasticity of substitution (Fontagné et al., 2022). The regressor here uses the rate in %, so the coefficient on RBT_{ijmt} is the elasticity of trade with respect to one plus the rate in %, which is comparable across

specifications but depends on the unit of the rate and is not itself an estimate of the elasticity of substitution.

The RBT retains a positive RVS under the EMT, IMT and PM structure, against exactly zero for the MFN rate under both structures that contain IMT effects (Table D4). The remaining variation is limited, because only 2.2% of observations carry a PSR. These observations come from 15 of the 26 importers in the sample. Ten of the 12 European Union member states among the importers supply none. No row carries one in 2022 or 2023 (Table D7), so the variation that identifies the elasticity comes from a small set of preferential cells. Multiplying the MFN rate by one minus the agreement indicator is not used, because it identifies a rate-by-agreement interaction and not a rate.

4. Results

This section proceeds in four steps. It first reports the main estimates and how they change with the FE structure, then compares the Poisson estimate with alternative estimators, then turns to differences across minerals and to estimates on a balanced block and finally examines how far the estimate can be relied on.

4.1 Main estimation

Table 1 reports PPML-HDFE estimates of exports on the RBT, AGR and four gravity covariates in eight FE specifications, with zeros retained. Columns (1) to (7) are the sequence conventionally reported in this literature, from additive exporter and importer effects in column (1) to exporter-year (EY), importer-year (IY), mineral-year and pair effects in column (7). Column (8) imposes the EMT, IMT and PM structure that the theory requires. Panel A reports the coefficients, the number of observations, the pseudo R^2 and the p-value of the regression equation specification error test (RESET) of the conditional mean, Panel B the RVS of each variable under the FE of the column and Panel C the FE. Covariates that vary only by pair are absorbed by the pair effects in columns (7) and (8).

As expected, the RBT coefficient is negative in every column, but its size depends strongly on the FE. Under additive country effects or EY and IY effects, in columns (1), (2) and (4), it lies between −1.57 and −2.52 and is significant at the 1% level. Once a mineral dimension enters, in columns (3), (5), (6) and (7), it falls to between −0.15 and −0.52 and is significant at the 5% level only in columns (5) and (6). In column (8) the estimate is −0.3257 and significant at the 1% level under pair-clustered SEs, a result that the reliability checks below qualify. It is close to the mineral-inclusive columns and between five and eight times smaller in absolute magnitude than in columns (1), (2) and (4).

Two features are consistent with the fall. Mineral-year effects remove shocks that move a mineral's trade and its preferences together. The sign depends on them (Panel F of Table E1). In addition, column (8) is the specification in which the RBT keeps the least variation, with an RVS of 0.0227 (Panel B), so the estimate rests on a small share of the variation in the rate. In this column the coefficient is identified from changes over time in a pair's RBT for a mineral, relative to the exporter's sales of that mineral to other importers and the importer's purchases of it from other exporters in the same year. The RESET does not reject column (8) (p=0.760) while it rejects column (4), which has no mineral dimension (p=0.001).

The remaining coefficients are as expected except that of AGR. In the six columns where they are identified DIST is negative (−0.81 to −0.93), within the range of −0.8 to −1 that Larch et al. (2021) report as usual for gravity models. CONTIG is positive (0.46 to 0.54). Both are significant at the 5% level or better. LANG and COL are never significant. AGR is unstable: it is positive and significant at the 10% level or better in columns (1), (3), (5) and (6), negative and significant in column (2) and close to zero in columns (7) and (8), so it depends on the FE as the tariff coefficient does.

Table 1. Main estimates by fixed-effect specification.

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Estimated coefficients								
RBT	-1.5738***	-1.7063***	-0.1490	-2.5213***	-0.4171**	-0.5179***	-0.2834	-0.3257***
	(0.1353)	(0.1158)	(0.1011)	(0.1726)	(0.1667)	(0.1655)	(0.1803)	(0.1147)
AGR	0.2598*	-0.2888**	0.2895**	-0.1529	0.3524**	0.3435**	0.0412	0.0002
	(0.1370)	(0.1445)	(0.1427)	(0.1634)	(0.1509)	(0.1501)	(0.1009)	(0.1061)
DIST	-0.8050***	-0.9296***	-0.8899***	-0.9139***	-0.8881***	-0.8863***	-	-
	(0.0803)	(0.0828)	(0.0897)	(0.0890)	(0.0934)	(0.0931)		
CONTIG	0.4797***	0.4834***	0.4556**	0.5375***	0.4678**	0.4682**	-	-
	(0.1793)	(0.1728)	(0.1950)	(0.1846)	(0.1971)	(0.1961)		
LANG	-0.1876	-0.1212	-0.1553	-0.1538	-0.1828	-0.1790	-	-
	(0.2660)	(0.2575)	(0.2388)	(0.2672)	(0.2319)	(0.2322)		
COL	0.0884	0.1489	0.3515	-0.0015	0.2889	0.2822	-	-
	(0.5291)	(0.5312)	(0.5591)	(0.5865)	(0.5717)	(0.5701)		
Observations	159,157	159,157	159,157	154,230	154,230	154,230	142,358	89,994
Pseudo R^2	0.5560	0.5937	0.7947	0.6576	0.8255	0.8301	0.9165	0.9675
RESET p-value	0.077	0.461	0.874	0.001	0.854	0.884	0.342	0.760
Panel B: Variance shares								
RBT	0.5250	0.4878	0.2863	0.4135	0.2121	0.2036	0.1938	0.0227
AGR	0.6324	0.5926	0.5926	0.4992	0.4992	0.4992	0.1216	0.1216
DIST	0.5409	0.5408	0.5408	0.5330	0.5330	0.5330	0.0000	0.0000
CONTIG	0.9183	0.9183	0.9183	0.9133	0.9133	0.9133	0.0000	0.0000
LANG	0.8052	0.8052	0.8052	0.7989	0.7989	0.7989	0.0000	0.0000
COL	0.8427	0.8427	0.8427	0.8390	0.8390	0.8390	0.0000	0.0000
Panel C: Fixed effects								
E and I	Yes	Yes	Yes	No	No	No	No	No
Y	No	Yes	Yes	No	No	No	No	No
M	No	No	Yes	No	Yes	No	No	No
EY and IY	No	No	No	Yes	Yes	Yes	Yes	No
MY	No	No	No	No	No	Yes	Yes	No
PAIR	No	No	No	No	No	No	Yes	No
EMT, IMT and PM	No	No	No	No	No	No	No	Yes

Notes: Significance levels: *** p<0.01, ** p<0.05, * p<0.1, SE in parentheses below each coefficient, SEs are pair-clustered, PPML-HDFE estimates of exports in thousands of US dollars, RBT is the log of one plus the realised bilateral tariff in %, a hyphen marks a covariate absorbed by the pair effects, RESET is the p-value of the test on the squared fitted linear index (Santos Silva & Tenreyro, 2022), Panel B reports the RVS under the FE of each column in the 159,157-row sample, E, I, Y and M denote exporter, importer, year and mineral effects, MY is a mineral-year effect, PAIR is an exporter-importer effect common to all minerals, AGR is the agreement indicator, DIST is the log of distance in km, CONTIG, LANG and COL are indicators of contiguity, common official language and common coloniser.

The pseudo R^2 rises from 0.56 in column (1) to 0.97 in column (8) as the FE absorb more of the variation in trade. The sample falls from 159,157 to 89,994 observations because cells are absorbed as singletons or separated once the three sets of effects are imposed together. The coefficient of column (8) is unchanged under a tighter solver tolerance and when the agreement indicator is dropped. It keeps its sign and significance when the rate enters in levels, when 2020 to 2023 are excluded and when estimation starts in 2005.

The size of the response also depends on how the rate enters, which Panel A of Table E1 shows for three forms: the logarithm of one plus the rate in % that Table 1 uses, the logarithm of one plus the rate as a fraction and the rate in levels. The fraction form is the one that the gravity system implies, because the tariff factor multiplies the delivered price and the coefficient is then the elasticity of trade with respect to the tariff factor. All three forms give a negative coefficient significant at the 1% level. The fraction and levels forms agree closely and

imply that a cut from 10% to 5% raises trade by about 114% to 123% and a cut from 5% to 0% by 122% to 123%, with wide intervals, whereas the form of Table 1 implies 22% and 79%. At the mean realised rate the elasticity is −0.43 for the fraction and levels forms against −0.24 for the form of Table 1. Table 1 keeps the latter for a common scale across its columns, so Panel A of Table E1 is the guide to magnitude.

For reference, Fontagné et al. (2022) estimate product-level tariff elasticities by PPML with importer-year and exporter-year effects and report an average trade elasticity, which is one plus the tariff coefficient, of −18.5 for the mineral products section of the HS and between −5.3 and −7.6 across all products depending on how positive and insignificant estimates are treated. The fraction-form coefficient here, −16.3, is a tariff coefficient and compares with about −19.5 implied by their section average, so it is of the same order. The designs differ, because their regressions are run product by product with no pair effect and their section does not coincide with the five minerals, so the comparison is indicative only. Boehm et al. (2023) find the same dependence on bilateral effects for tariffs on all goods, where log-levels estimates of −3.7 to −7.0 fall to about −1 once pair-product effects are included. Esteve-Pérez et al. (2020) report the same for the General Agreement on Tariffs and Trade and World Trade Organization indicator, where a positive effect of 0.365 to 0.480 without pair effects becomes −0.027 once exporter-time, importer-time and pair effects are all included.

Pooling minerals in the effect structure keeps the negative sign wherever a mineral-year control is retained (Panel F of Table E1). With EMT and IMT effects and one pair effect common to all minerals the coefficient is −0.2813, close to column (8). With EY, IY and mineral-year effects and PM effects it is −0.2236. Both are significant at the 5% level. Dropping the mineral-year effect gives 0.0204, so the sign depends on the control for mineral-specific shocks. The MFN rate keeps an RVS of 0.0320 only under the second structure, because the importer effect no longer absorbs differences across minerals. Its coefficient (−0.1681) is not significant.

4.2 Alternative estimators

Table 2 sets the Poisson estimate beside the alternative estimators that can be run on the same regressor and sample. Columns (1) to (4) use EY and IY effects and columns (5) to (8) the EMT, IMT and PM structure, so every estimator appears under a structure that is common in the literature and under the structure that the theory requires. The estimators are OLS-L1P, OLS-LN, GPML, PPML-HDFE and the bias-corrected PPML-BC. The OLS columns are the within estimator of Balázsi et al. (2018) for the two transformations of exports.

Table 2. Alternative estimators applied to the same data.

	Exporter-year and importer-year effects				**EMT, IMT and PM effects**			
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
Variable	**OLS-L1P**	**OLS-LN**	**GPML**	**PPML**	**OLS-L1P**	**OLS-LN**	**PPML**	**PPML-BC**
Panel A: Estimated coefficients								
RBT	-1.2426***	-2.1171***	-2.3783***	-2.5213***	0.0535	-0.0730	-0.3257***	-0.3111***
	(0.0690)	(0.1146)	(0.0835)	(0.1726)	(0.0607)	(0.1076)	(0.1147)	(0.1147)
AGR	-0.0275	-0.1065	-0.0076	-0.1529	0.1467***	0.0916	0.0002	0.0154
	(0.0855)	(0.1364)	(0.1090)	(0.1634)	(0.0497)	(0.1063)	(0.1061)	(0.1061)
DIST	-0.9087***	-1.0138***	-0.7698***	-0.9139***	-	-	-	-
	(0.0652)	(0.0650)	(0.0591)	(0.0890)				
CONTIG	0.3871**	0.1391	0.0507	0.5375***	-	-	-	-
	(0.1696)	(0.1824)	(0.1465)	(0.1846)				
LANG	0.2416**	-0.0238	-0.1226	-0.1538	-	-	-	-
	(0.1068)	(0.1685)	(0.1717)	(0.2672)				

COL	0.3751 (0.2423)	0.4651 (0.3081)	0.0322 (0.3004)	-0.0015 (0.5865)	-	-	-	-
Observations	159,157	48,242	48,278	154,230	159,157	47,108	89,994	89,994
R^2 or pseudo R^2	0.3228	0.3275	-	0.6576	0.7978	0.8048	0.9675	0.9675

Notes: Significance levels: *** p<0.01, ** p<0.05, * p<0.1, SE in parentheses below each coefficient, SEs are pair-clustered, column (4) is column (4) of Table 1 and column (7) is column (8) of Table 1, PPML-BC applies the correction of Weidner and Zylkin (2021) to column (7) with the SEs of the uncorrected fit, GPML did not converge under the EMT, IMT and PM structure, a hyphen marks a covariate absorbed by the pair effects or a statistic that is not defined, OLS-LN and GPML drop zero flows, AGR, DIST, CONTIG, LANG and COL are defined in Table 1.

Under EY and IY effects the four estimators agree that the RBT coefficient is negative and significant at the 1% level, from −1.24 for OLS-L1P to −2.52 for PPML and that DIST is negative and significant at the 1% level, between −0.77 and −1.01. Under the required structure the estimators separate. PPML returns −0.3257 and PPML-BC −0.3111, both significant at the 1% level, whereas OLS-L1P returns 0.0535 and OLS-LN −0.0730 and neither differs from zero. The log-linear estimators therefore lose the response once the required structure is imposed.

The bias correction moves the RBT coefficient by only 0.0146, so the incidental-parameter bias that Weidner and Zylkin (2021) characterise does not change the conclusion at this sample size. In Table 2 the log-linear estimators shrink sharply when the EMT, IMT and PM structure replaces the EY and IY structure, from −1.2426 to 0.0535 for OLS-L1P and from −2.1171 to −0.0730 for OLS-LN, as PPML-HDFE does. A divergence between log-linear and Poisson estimates arises in the three-way FE gravity model of Larch et al. (2019), who trace it to heteroscedasticity in the log-linear residuals. The samples differ across estimators, because OLS-LN discards zero trade, which makes up most of the sample, whereas PPML-HDFE loses observations to separation, so Table 2 does not separate the effect of the estimator from the effect of the sample.

4.3 Heterogeneity and estimates on a balanced block

Estimating column (8) separately by mineral shows that the pooled result hides differences in sign (Panel A of Table 3). Farag and Zaki (2025) likewise find that the response of critical mineral trade to political distance differs across individual minerals. The elasticity is negative and significant for cobalt, copper and graphite, positive and significant for lithium and not different from zero for the rare earths. The incidental-parameter correction in Panel B changes no mineral elasticity by more than 0.07 and leaves the sign and significance of each as they are. The pooled estimate is close to the copper elasticity (−0.3257 and −0.3525) because copper carries 93% of the information on the rate. A Wald test in a pooled model with mineral-specific slopes rejects equal slopes across the five minerals (chi-squared 31.87 on 4 degrees of freedom, p<0.001) but not across the four minerals other than lithium (4.11 on 3 degrees of freedom, p=0.25), so the heterogeneity is a difference between lithium and the rest. Each mineral rests on fewer PSR rows than the pooled sample, so the mineral-level estimates are indicative.

The positive lithium estimate is not explained. Lithium has about as many identifying rows as the other minerals, so limited identifying variation does not account for it. The conventional sequence reports one number for all five minerals and nothing in columns (1) to (7) of Table 1 would reveal that the elasticities differ in sign. Slope constancy tests for three-dimensional panels (Jin et al., 2025) examine whether slopes vary over exporters or importers and not across a fourth dimension, so equality is tested across minerals instead.

Beyond the linear Poisson specification, the nonparametric estimator relaxes the assumption that the response is linear in the RBT and the interactive-effects estimators relax the assumption that additive effects remove all common variation. Both need a complete array, so Panels C to G of Table 3 use the largest complete block of the estimation sample, with 28

exporters that do not appear as importers, 19 importers and the 24 years from 2000 to 2023 in each mineral, which gives 63,840 rows of which 84% record no trade. The block is not the sample of column (8), because copper supplies 96% of its trade value, 2.6% of its rows carry a PSR and the observations that PPML drops leave between 2,352 and 9,672 rows for each mineral. Panel C reports the Poisson estimates on the block and Panels D to G the OLS-L1P, nonparametric, 3DCCE and multi-level IFE estimates, which use the log of one plus exports as the dependent variable because they need an additive mean.

Table 3. Estimates by mineral and on the balanced block.

Variable	(1) Pooled	(2) Cobalt	(3) Copper	(4) Graphite	(5) Lithium	(6) Rare earths
Panel A: PPML, full sample						
RBT	-0.3257*** (0.1147)	-0.7090*** (0.2117)	-0.3525*** (0.1215)	-0.4404*** (0.1324)	0.7124*** (0.2082)	-0.0071 (0.3370)
AGR	0.0002 (0.1061)	-0.1292 (0.2676)	-0.0076 (0.1086)	0.2702 (0.2170)	0.4702 (0.3640)	-0.1259 (0.4577)
Observations	89,994	14,343	26,687	20,691	13,228	15,045
Pseudo R^2	0.9675	0.9697	0.9552	0.9482	0.9879	0.9743
Information weight	-	0.007	0.929	0.030	0.029	0.006
Panel B: Bias-corrected PPML						
RBT	-0.3111*** (0.1147)	-0.7391*** (0.2117)	-0.3379*** (0.1215)	-0.5080*** (0.1324)	0.7683*** (0.2082)	-0.0209 (0.3370)
AGR	0.0154 (0.1061)	-0.0395 (0.2676)	0.0097 (0.1086)	0.2741 (0.2170)	0.4586 (0.3640)	-0.2648 (0.4577)
Panel C: PPML, balanced block						
RBT	-0.1408 (0.1212)	3.7051** (1.6315)	-0.1768 (0.1202)	0.6034* (0.3567)	0.8860*** (0.2570)	-1.9649 (1.9332)
AGR	-0.1799* (0.1048)	-2.7042 (2.1717)	-0.1799* (0.1018)	-0.4514 (0.7235)	1.0154*** (0.3662)	-0.4333 (0.9529)
Observations	22,690	2,434	9,672	5,238	2,352	2,994
Pseudo R^2	0.9805	0.9924	0.9743	0.9228	0.9941	0.9892
Panel D: OLS-L1P, balanced block						
RBT	0.1022 (0.0674)	0.5183** (0.2080)	-0.0087 (0.3451)	-0.0410 (0.1062)	0.0808 (0.1004)	0.0905 (0.1082)
AGR	0.0157 (0.0761)	-0.1237 (0.1093)	-0.1405 (0.2742)	-0.0297 (0.1022)	0.0880 (0.0585)	0.2851* (0.1668)
Observations	63,840	12,768	12,768	12,768	12,768	12,768
R^2	0.7986	0.5943	0.7735	0.7345	0.8441	0.5740
Panel E: Nonparametric, balanced block						
Average derivative of RBT	0.0878 (0.0821)	0.5105* (0.2617)	-0.2114 (0.4419)	-0.0349 (0.0991)	0.0639 (0.0920)	0.0897 (0.1059)
AGR	0.0155 (0.0702)	-0.1251 (0.1058)	-0.1322 (0.2767)	-0.0297 (0.0921)	0.0880 (0.0541)	0.2850* (0.1572)
Gradient at a rate of 0%	0.1546 (0.1637)	0.6953* (0.3696)	-0.2475 (0.5007)	-0.0110 (0.1138)	0.0472 (0.0812)	0.1063 (0.1211)
Gradient at a rate of 6%	0.0308 (0.1290)	-	-	-0.0490 (0.0992)	0.0634 (0.1025)	0.0767 (0.1079)
Bandwidth	1.345	0.851	1.825	2.883	2.883	2.883
Panel F: Pooled 3DCCE						
RBT	0.0866 (0.1607)	0.5714* (0.3449)	0.6418 (0.5226)	-0.0208 (0.2132)	-0.3587 (0.3937)	-0.0407 (0.1766)
AGR	-0.1980 (0.1972)	-0.1984 (0.3321)	-0.5322 (0.7573)	0.2027 (0.3354)	-0.0562 (0.2424)	-0.3924 (0.2490)
Share of pair series with variation	0.422	0.445	0.445	0.404	0.406	0.408

Panel G: Multi-level IFE						
RBT	-0.0053 (0.0185)	-0.0379 (0.0459)	-0.0657 (0.1353)	-0.0096 (0.0688)	0.0776* (0.0466)	0.0334** (0.0136)
AGR	-0.0956* (0.0508)	-0.1554** (0.0729)	-0.2823 (0.1952)	-0.0083 (0.0538)	0.0267 (0.0368)	-0.0452 (0.0420)
Global factors	1 to 2	1	2	1	1	2

Notes: Significance levels: *** p<0.01, ** p<0.05, * p<0.1, SE in parentheses below each coefficient, SEs are pair-clustered in Panels A to D and allow for dependence within each exporter and importer in Panels F and G, Panel A is column (8) of Table 1 by mineral, the information weight is each mineral's share of the fitted mean times the squared residualised rate, Panels C to G use the balanced block of 63,840 rows and Panels D to G the log of one plus exports, in Panel E SEs are from a pair bootstrap with 400 draws (300 per mineral) and stars use a normal approximation, the average derivative is the gradient averaged between the 5th and 95th percentiles of the RBT, a hyphen marks a rate above the 95th percentile, the share in Panel F is that of pair series in which a regressor varies after the projection on the averages, the pooled entry for global factors in Panel G is the range across minerals, AGR is defined in Table 1.

On the block none of the three estimators of the pooled slope is significant and their signs differ: the Poisson estimate is −0.1408, the OLS-L1P estimate 0.1022 and the nonparametric average derivative 0.0878. The estimated gradient of the pooled sample is positive at every rate shown and falls from 0.1546 at a rate of 0% to 0.0308 at a rate of 6%, but no gradient is significant. Only cobalt returns a significant OLS-L1P coefficient and a significant average derivative, both positive and the derivative at the 10% level. The Poisson estimates for cobalt and graphite are positive and significant at the 10% level or better, the opposite sign to Panel A.

Panels F and G add the two interactive-effects estimators. The pooled 3DCCE estimate of the RBT is 0.0866 and the pooled multi-level IFE estimate −0.0053, neither significant, so the pooled OLS-L1P estimate of 0.1022 falls to a value close to zero once the factors are removed. By mineral the RBT coefficient is positive and significant for lithium at the 10% level and for the rare earths at the 5% level under multi-level IFE and for cobalt at the 10% level under 3DCCE. It is not significant elsewhere. Both estimators rest on a linear model of an outcome that is zero in 84% of the block.

The block therefore neither reproduces nor refutes the full-sample result. It identifies the response from the 2.6% of rows that carry a PSR, where copper accounts for 99% of the trade value. No estimator returns a significant pooled response on the block, including the Poisson estimator, unlike column (8). What the evidence establishes is that the sign of the estimated response is not common to the estimators that can be applied to a gravity panel of this form, so the negative response in Table 1 is a property of the Poisson fit to the full sample and should be read with the reliability results that follow. The evidence cannot show which estimator is closer to the truth, because the linear estimators discard the multiplicative structure that motivates the Poisson fit.

4.4 Reliability of the estimate

How far the estimate can be relied on depends on four sets of checks. They concern the source of the identifying variation, stability across subperiods and countries, anticipation and selection and inference. Table E1 in Appendix E collects the numerical results of the simulation, the restricted samples, the pair bootstrap and the effective clusters.

First, the variation comes from a narrow set of rows. The rows whose RBT differs from the MFN rate come from 14 of the 26 importers in the sample and none falls in 2022 or 2023, so the estimate speaks to the importers and years that grant preferences and not to the sample as a whole. In the sample of column (8) these rows come from only 95 directed pairs.

Second, the estimate is not recovered within subperiods but does not depend on any single country. The three subperiod estimates, each for a window of ten years or less, are small, positive and not significant, perhaps because the PM effects absorb most of the variation

inside a short window, so the estimate should not be read as stable over time. Dropping each exporter and each importer in turn leaves the elasticity negative in every fit and significant in all but one, so no single country drives the sign.

Third, future rates predict current trade. Adding the rate one and two years ahead gives negative and significant coefficients on the leads and lowers the coefficient on the current rate. This is consistent with anticipation of preferential access, with selection into agreements and with a slow-moving rate whose current and future values proxy for the same variation. The data cannot separate them. An agreement signed because trade in a mineral was already rising would overstate the response and AGR does not absorb that selection. The estimate is therefore read as an association between preferential tariff cuts and trade.

Fourth, the significance of the estimate is sensitive to how uncertainty is computed. In a Monte Carlo simulation calibrated to the application PPML-HDFE is centred on the true value but the 95% interval based on pair-clustered SEs covers it in 87% to 92% of draws, below the nominal level at every preferential share examined (Panel B of Table E1). The pair bootstrap interval, from −0.64 to 0.01, includes zero (Panel D) and the effective number of clusters is far below the number of clusters at every level of clustering (Panel E of Table E1).

Restricting the sample to importer-mineral-years with at least one PSR or to importers that ever report one keeps all rows that differ from the MFN rate and the sign of the estimate, but the estimate is smaller in absolute magnitude and not significant at the 5% level in either sample (Panel C). If unrecorded preferences were attenuating the full-sample estimate the restriction would raise its magnitude and it does not. The data do not show why it falls. In sum, the Poisson estimate is negative and of the order of −0.3 in the specification that the theory requires, but its sign is not common to the estimators and subperiods examined and its significance depends on how uncertainty is computed.

5. Conclusion

This paper shows within the structural gravity system that a tariff schedule published at the importer and product level is not a bilateral trade cost. An origin-invariant rate cancels from the gravity equation and its RVS under IMT effects is zero, so treating it as a bilateral regressor produces coefficients that cannot be interpreted as elasticities. Using the RBT, which differs across origins through preferential access, PPML-HDFE with EMT, IMT and PM effects yields a negative tariff elasticity for critical minerals that is much smaller in absolute magnitude than the additive-effects estimates that dominate reporting, similar in size to specifications that include a mineral dimension and heterogeneous across minerals. Log-linear, nonparametric and interactive-effects estimators with the same effects do not reproduce a significant pooled response on the same data or on a balanced block of it, whereas an incidental-parameter bias correction leaves the Poisson estimate nearly unchanged and effect structures that pool minerals keep the sign where a mineral-year control is retained. The evidence base is narrow, because few observations carry a PSR, so the estimate is an elasticity identified from preferential cells and not an economy-wide response.

Three implications follow for policy appraisal. First, the evidence does not support a single claim that tariff reductions raise or reduce trade in critical minerals as a group, because the estimated response is negative for cobalt, copper and graphite, positive for lithium and absent for the rare earths, so appraisal should be mineral specific. Second, an MFN schedule cannot identify a bilateral response in this framework, so assessing a change in an MFN schedule requires origin-specific rates and fuller recording of partner-specific rates in tariff databases would widen the identifying variation available to future studies. Third, because the estimate is identified from preferential access and is an association, it should not be used to simulate the effect of an unanticipated tariff change without further identifying assumptions.

Several limitations bound the result. Observations without a PSR are assigned the MFN rate, so differences in preference margins that are not recorded are measured with error. WITS partner-group codes could not be expanded to countries and the final year and several

importer-year groups could not be retrieved, so the RVS is a lower bound. The identifying specification absorbs a large part of the sample as singletons or separated observations, which limits precision. The estimate is not recovered within short subperiods, future rates predict current trade, its simulated interval undercovers and it is not significant at the 5% level in the restricted samples, under a pair bootstrap or with the log-linear, nonparametric and interactive-effects estimators. The RVS bounds identifying variation without bounding bias, so a regressor with a large RVS may still be endogenous. The wider point is that the FE structures now standard in gravity work can remove a regressor entirely and that a regressor left with almost no variation passes unnoticed. Reporting what each specification leaves behind adds one row to a results table.

Data availability statement

All data are publicly available from the sources named in the manuscript. The estimation panel, code and stored results are deposited at Zenodo https://doi.org/10.5281/zenodo.23400587

Statements and Declarations

Funding. This research received no external funding.

Competing interests. The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics approval and consent. Not applicable. The study uses public aggregate trade and tariff data.

Clinical trial number. Not applicable.

Use of artificial intelligence tools. An artificial intelligence assistant (Claude, Anthropic) was used to write and debug estimation code, to run the analyses and to draft parts of the text. Grammarly was used for language editing. The author specified the research design, checked the results and the references and takes full responsibility for the content.

Appendix A. Structural system, estimator and variance

A.1 The structural system

For exporter i, importer j, mineral m and year t the structural system consists of the gravity equation, the two multilateral resistance terms and market clearing, with every sum running over all partners in the model, domestic sales included:

$$X_{ijmt} = \frac{Y_{imt}E_{jmt}}{Y_{mt}}\left(\frac{\tau_{ijmt}}{\Pi_{imt}P_{jmt}}\right)^{1-\sigma_m}, \quad \tau_{ijmt} = \left(1+\theta_{ijmt}\right)\kappa_{ijmt}$$

$$\Pi_{imt}^{1-\sigma_m} = \sum_j \left(\frac{\tau_{ijmt}}{P_{jmt}}\right)^{1-\sigma_m}\frac{E_{jmt}}{Y_{mt}}, \quad P_{jmt}^{1-\sigma_m} = \sum_i \left(\frac{\tau_{ijmt}}{\Pi_{imt}}\right)^{1-\sigma_m}\frac{Y_{imt}}{Y_{mt}}$$

$$Y_{imt} = \sum_j X_{ijmt}, \quad E_{jmt} = \sum_i X_{ijmt}, \quad Y_{mt} = \sum_i Y_{imt} = \sum_j E_{jmt}$$

Taking logarithms of the gravity equation and collecting the terms that depend on one side of the market only defines the structural effects,

$$\alpha_{imt} = \ln Y_{imt} - (1-\sigma_m)\ln \Pi_{imt}, \quad \gamma_{jmt} = \ln E_{jmt} - \ln Y_{mt} - (1-\sigma_m)\ln P_{jmt}$$

so that $\ln X_{ijmt} = \alpha_{imt} + \gamma_{jmt} + (1-\sigma_m)\ln\tau_{ijmt}$, the delivered value of trade. Writing the bilateral frictions as a pair-mineral component η_{ijm} plus a remainder and collecting the policy variables in x_{ijmt} with parameters β, which equal $-\sigma_m$ for the log of one plus the tariff expressed as a fraction because BACI values are FOB (Fontagné et al., 2022), gives the estimating equation with a multiplicative error whose conditional mean is one:

$$X_{ijmt} = \exp\left(\alpha_{imt} + \gamma_{jmt} + \eta_{ijm} + x_{ijmt}'\beta\right)u_{ijmt}, \quad \mathbb{E}\left[u_{ijmt} \mid \alpha_{imt}, \gamma_{jmt}, \eta_{ijm}, x_{ijmt}\right] = 1$$

A.2 Design matrices and the residual maker

Order the complete array by exporter, importer, mineral and year, with N_i, N_j, N_m and N_t the numbers of exporters, importers, minerals and years, I_N an identity matrix and ι_N a vector of ones, in the notation of Balázsi et al. (2018). The three sets of effects of the preferred specification, exporter-mineral-year, importer-mineral-year and pair-mineral, have the blocks D_{emt}, D_{imt} and D_{pm}, stacked in D_F, of the design:

$$D_{emt} = I_{N_i} \otimes \iota_{N_j} \otimes I_{N_m N_t}, \quad D_{imt} = \iota_{N_i} \otimes I_{N_j N_m N_t}, \quad D_{pm} = I_{N_i N_j N_m} \otimes \iota_{N_t}$$

Each block has as many rows as the array has cells, $N_c = N_i N_j N_m N_t$. The sets of effects overlap, because a function of exporter, mineral and year that is also a function of importer, mineral and year can depend on mineral and year alone. The same holds for the other two pairs and for all three together. Counting the overlaps by inclusion and exclusion gives the rank of the design for a complete array, which equals N_c minus the rank of the residual maker:

$$\operatorname{rank}(D_F) = N_c - (N_i - 1)\big(N_j - 1\big)N_m(N_t - 1)$$

The residual maker for a complete array has the closed form in the first display below, which holds because its range is the set of arrays whose sums over exporter, importer and year vanish in every cell of the remaining indices. The estimation sample is a row selection of the array, because own-country flows are absent and separated cells are dropped. With N_c the size of the complete array, n the number of retained rows and L the $n \times N_c$ matrix of the rows of the identity matrix that correspond to retained observations, the projection computed on the rows actually used is the second display, following the selection-matrix treatment of self-flows in Balázsi et al. (2018):

$$\mathbf{M}_F = Q_{N_i} \otimes Q_{N_j} \otimes I_{N_m} \otimes Q_{N_t}, \quad Q_N = I_N - N^{-1}\iota_N \iota_N{}'$$
$$\mathbf{M}_{F,L} = I_n - LD_F(D_F{}'L'LD_F)^{-}D_F{}'L'$$

Multiplying out the Kronecker product gives eight terms, which are, term by term, the scalar transformation below, in which a dot replaces an index that is averaged over and the sign alternates with the number of averaged indices:

$$\ddot{x}_{ijmt} = x_{ijmt} - \bar{x}_{\cdot jmt} - \bar{x}_{i\cdot mt} - \bar{x}_{ij\cdot m} + \bar{x}_{\cdot\cdot mt} + \bar{x}_{\cdot j\cdot m} + \bar{x}_{i\cdot\cdot m} - \bar{x}_{\cdot\cdot\cdot m}$$

The complete-array form is exact only when every cell is observed. On the estimation sample the three sets of effects are removed one at a time, with $B_1 = LD_{\mathrm{emt}}$, $B_2 = LD_{\mathrm{imt}}$ and $B_3 = LD_{\mathrm{pm}}$, in the manner of the stepwise projections of Balázsi et al. (2018). After the last step the result equals $\mathbf{M}_{F,L}$ above, because projecting in sequence onto orthogonal complements removes the span of all three designs:

$$\mathbf{M}^{(k)} = \mathbf{M}^{(k-1)} - \mathbf{M}^{(k-1)}B_k\big(B_k{}'\mathbf{M}^{(k-1)}B_k\big)^{-}B_k{}'\mathbf{M}^{(k-1)}, \quad k = 1,2,3, \quad \mathbf{M}^{(0)} = I_n$$

A.3 The PPML estimator

Let Ω index the retained observations, x_o the policy regressors and d_o the row of the design for observation o, with δ stacking all effects. The Poisson pseudo-log-likelihood is:

$$\ell(\beta, \delta) = \textstyle\sum_{o\in\Omega}\left[X_o \ln \mu_o - \mu_o\right], \quad \mu_o = \exp(x_o{}'\beta + d_o{}'\delta)$$

The first-order conditions are the first display below, in which the second condition states that fitted and observed totals agree within every effect cell, the property that makes the multilateral resistance terms recoverable exactly (Fally, 2015). Concentrating out the effects by weighted projection gives the residualised regressor and the concentrated score for the policy parameters in the second display, where the cell totals make the projection term vanish, so that the concentrated score equals the unconcentrated score:

$$\textstyle\sum_{o\in\Omega} x_o\,(X_o - \mu_o) = 0, \quad \sum_{o\in\Omega} d_o\,(X_o - \mu_o) = 0$$
$$\ddot{x} = x - D\big(D'\widehat{W}D\big)^{-}D'\widehat{W}x, \quad \widehat{W} = \operatorname{diag}(\hat{\mu}_o), \quad \textstyle\sum_{o\in\Omega}\ddot{x}_o\,(X_o - \hat{\mu}_o) = 0$$

The estimator iterates the profile Newton step, each iteration being a weighted within transformation of the kind that routines for high-dimensional fixed effects deliver (Correia et al., 2020), where $\hat{H}$ is the negative of the profile Hessian:

$$\beta^{(k+1)} = \beta^{(k)} + \hat{H}^{-1} \sum_{o\in\Omega} \ddot{x}_o \, (X_o - \hat{\mu}_o), \quad \hat{H} = \sum_{o\in\Omega} \hat{\mu}_o \, \ddot{x}_o \ddot{x}_o{}'$$

The coefficient is consistent for the conditional mean under heteroskedasticity of unknown form (Santos Silva & Tenreyro, 2006). For Poisson with strictly exogenous regressors the leading incidental-parameter terms in the panel setting of Fernández-Val and Weidner (2016, p. 299) vanish and the three-way gravity case is characterised by Weidner and Zylkin (2021). Within a mineral the structure is the three-way case and the pooled estimator combines five such panels under one slope. A bias expansion for that common slope is not given in the papers consulted for this study, so the headline estimates are reported without bias correction and the stacked correction is reported as a check.

A.4 Inference

Standard errors use the cluster-robust sandwich in which the directed pair defines the cluster c, with G clusters, so that serial dependence within a pair is left unrestricted:

$$\hat{V} = \hat{H}^{-1}\big(\textstyle\sum_{c=1}^{G} s_c \, s_c{}'\big)\hat{H}^{-1}, \quad s_c = \textstyle\sum_{o\in c} \ddot{x}_o \, (X_o - \hat{\mu}_o)$$

A leave-one-cluster-out jackknife addresses the downward bias of the sandwich (MacKinnon et al., 2023; Pfaffermayr, 2026) but was not computed for the full design, because demeaning failed in the cluster refits, so the reported standard errors are those of the sandwich up to any finite-sample multiplier applied by the estimation software. The theory of the feasible bias-corrected version is not yet established (Pfaffermayr, 2026).

Appendix B. Cancellation of an origin-invariant trade cost

Write the bilateral trade cost as $\tau_{ijmt} = T_{jmt}\,\kappa_{ijmt}$, where T_{jmt} varies only with the importer, the mineral and the year and κ_{ijmt} collects the remaining bilateral frictions. The inward resistance term and the gravity equation take the forms in the main text and the outward resistance term becomes:

$$\Pi_{imt}^{1-\sigma_m} = \textstyle\sum_j \left(\frac{T_{jmt}\kappa_{ijmt}}{P_{jmt}}\right)^{1-\sigma_m} \frac{E_{jmt}}{Y_{mt}} = \sum_j \frac{\kappa_{ijmt}^{1-\sigma_m}}{A_{jmt}} \frac{E_{jmt}}{Y_{mt}}$$

so that T_{jmt} drops from the outward resistance term as well as from the gravity equation. In logarithms, with the structural effects of Appendix A,

$$(1-\sigma_m)\ln P_{jmt} = (1-\sigma_m)\ln T_{jmt} + \ln A_{jmt}$$
$$\ln X_{ijmt} = \alpha_{imt} + \left[\ln E_{jmt} - \ln Y_{mt} - \ln A_{jmt}\right] + (1-\sigma_m)\ln \kappa_{ijmt}$$

the component $(1-\sigma_m)\ln T_{jmt}$ in the trade cost is offset exactly by the same term inside the inward resistance and the tariff enters only through the aggregates E_{jmt} and Y_{imt}, which the importer-mineral-year and exporter-mineral-year effects absorb. The cancellation requires that T_{jmt} be common to every origin i in the sum defining P_{jmt}, including the domestic origin when domestic sales are in the system. Trade values in BACI are FOB, so observed flows carry one further factor, the inverse of the tariff factor ${T_{jmt}}^{-1}$, which is absorbed by the importer-mineral-year effect, so the conclusion is unchanged.

The moment function of Jochmans (2017) compares the products of outcomes in two arrangements of the same four cells, $y_{ij}\,y_{i'j'}$ and $y_{ij'}\,y_{i'j}$, each divided by the corresponding product of means $\varphi_{ij}\,\varphi_{i'j'}$ and $\varphi_{ij'}\,\varphi_{i'j}$, where $\varphi_{ij} = \exp\big(x_{ij}{}'\psi\big)$. For a regressor that depends on the importer alone, $x_{ij} = b_j$, the two products of means coincide for every value of the parameter,

$$\varphi_{ij}\varphi_{i'j'} = \exp\big(b_j\psi + b_{j'}\psi\big) = \varphi_{ij'}\varphi_{i'j}$$

so the two terms of the moment function differ only in the arrangement of the outcomes and not in the weights attached to them and the population moment is zero at every value of ψ. The same holds for a regressor that depends on the exporter alone. The result holds for any instrument function of the four regressor values by which Jochmans (2017) multiplies the moment function. It is a consequence of the form of the moment function worked out here and not a statement made by Jochmans (2017).

The moment function of Yang and Zhang (2023) compares the products of outcomes in two arrangements of eight cells, $y_{ijt}\,y_{i\prime j\prime t}\,y_{ij\prime s}\,y_{i\prime js}$ and $y_{ijs}\,y_{i\prime j\prime s}\,y_{ij\prime t}\,y_{i\prime jt}$, each divided by the corresponding product of means. For a regressor that depends on the importer and year alone, $x_{ijt} = b_{jt}$, the two products of means coincide for every value of the parameter, because each contains the four values b_{jt}, $b_{j\prime t}$, b_{js} and $b_{j\prime s}$ once. The same holds for a regressor that depends on the exporter and year alone or on the pair alone, so the population moment is zero at every value of the parameter for any instrument function of the eight regressor values. The signed sum over the eight cells of the derivative of the log mean function, which multiplies the instrument function in the Jacobian of Yang and Zhang (2023), vanishes identically in the same cases, so their full-rank condition fails whatever the instrument function. These extensions are the author's own algebra and not statements made by Yang and Zhang (2023), who note a related failure of the rank condition for a particular instrument function when the temporal difference of a regressor is a deterministic function of the two years.

Appendix C. The residual variance share

Let $\tilde{x}$ denote the regressor in deviation from its overall mean. Computed on the rows actually used, the residual variance share is:

$$\rho(x,F) = \frac{\tilde{x}'\mathbf{M}_{F,L}\tilde{x}}{\tilde{x}'\tilde{x}} = 1 - R^2(\tilde{x} \text{ on } LD_F), \quad \tilde{x} = x - \bar{x}\,\iota_n$$

For a complete array it is computed with the Kronecker form of the residual maker in Appendix A. The profiles below show that a regressor that is a function of a subset of the indices is annihilated by a structure that contains that subset, with h a vector of values indexed by the subset. The applied tariff and destination concentration are removed entirely while the agreement indicator remains, because no effect in the structure is indexed on a subset containing both the pair and the year:

$$x = \left(\iota_{N_i} \otimes I_{N_j N_m N_t}\right) h \;\Rightarrow\; \mathbf{M}_F x = 0, \quad \text{importer-mineral-year profile}$$

$$x = \left(I_{N_i} \otimes \iota_{N_j} \otimes I_{N_m N_t}\right) h \;\Rightarrow\; \mathbf{M}_F x = 0, \quad \text{exporter-mineral-year profile}$$

$$x = \left(I_{N_i N_j} \otimes \iota_{N_m} \otimes I_{N_t}\right) h \;\Rightarrow\; \mathbf{M}_F x \neq 0, \quad \text{pair-year profile}$$

The share is computed without forming the dummy matrix. Each set of effects defines a group-mean operator $\mathcal{P}_f$, with $q_f(o)$ the cell of observation o under the subset f. Applying the complements of these operators in turn converges to the projection used in the share (Correia et al., 2020):

$$\left(\mathcal{P}_f z\right)_o = \frac{1}{|q_f(o)|}\sum_{o\prime \in q_f(o)} z_{o\prime}, \quad \mathbf{M}_{F,L}\tilde{x} = \lim_{k\to\infty}\left[\left(I - \mathcal{P}_{f_K}\right)\cdots\left(I - \mathcal{P}_{f_1}\right)\right]^k \tilde{x}$$

Because the variation available to identify a parameter can only shrink as effects are added, a larger structure never raises the share. If the variation profile $\mathcal{S}(x)$ of a regressor lies within a subset f that belongs to the required structure F but not to the candidate set $\mathcal{C}$ of a selection procedure, the share is zero although the procedure never evaluates the effect that would absorb it:

$$F \subseteq F' \Rightarrow \mathbf{M}_{F'}\mathbf{M}_F = \mathbf{M}_{F'} \Rightarrow \rho(x,F') \leq \rho(x,F), \quad \mathcal{S}(x) \subseteq f,\; f \in F\backslash\mathcal{C} \Rightarrow \rho(x,F) = 0$$

Appendix D. Preliminary checks, descriptive statistics and product lines

Table D1 traces the construction of the samples and Table D2 defines the variables. Table D3 reports their descriptive statistics and Table D4 the RVS of the policy regressors under four nested FE structures. Table D5 lists the product lines attributed to each mineral, Table D6 the importers and exporters and Table D7 the rows with a PSR in each year. Trade is zero in 69.7% of the rows of sample E and a PSR is retrieved for 2.2% of them, of which 74.8% carry a rate that differs from MFN.

Table D1. Construction of the samples.

Sample	Rule	Rows	Share of zero flows
Panel A: Sample rules and sizes			
Complete panel	50 exporters, 50 importers, 2,464 directed pairs, 5 minerals, 1995 to 2024	369,600	0.7675
Sample A	MFN-G, AGR and DIST observed	186,383	0.7018
Sample E	RBT, AGR, DIST and MFN-G observed, 49 exporters, 26 importers, 1,254 directed pairs, 1995 to 2023	159,157	0.6967
Balanced block	28 exporters, 19 importers, 24 years, 5 minerals, every cell observed in sample E	63,840	0.8440
Panel B: Rows retained by PPML-HDFE			
Sample	**Fixed-effect structure**	**Rows retained**	**Share of sample**
Sample E	(1) EY and IY	154,230	0.9690
Sample E	(2) EY, IY and PAIR	142,358	0.8945
Sample E	(3) EMT and IMT	111,201	0.6987
Sample E	(4) EMT, IMT and PM	89,994	0.5654
Balanced block	(4) EMT, IMT and PM	22,690	0.3554

Notes: Sample A keeps the cells of the complete panel on which MFN-G, AGR and DIST are observed, sample E keeps the cells of sample A on which the RBT is also observed, the balanced block keeps the 19 importers, the 28 exporters that are not importers in sample E and the 24 years for which every mineral is observed in sample E, the share of zero flows is the share of cells with TRADE equal to zero, rows retained by PPML-HDFE exclude singleton and separated observations, the share of the sample is the rows retained divided by the rows of the sample, structures (1) to (4) are defined in Table D4 and variables in Table D2, the complete panel includes S19, the BACI code for Other Asia, not elsewhere specified, which has no distance or agreement data and so is in neither sample A nor sample E, dropping the MFN-G condition adds 143 rows to sample E and leaves the estimate of column (8) of Table 1 unchanged.

Table D2. Variable definitions.

Label	Definition	Source	Unit
Panel A: Trade and tariffs			
TRADE	Bilateral export value of the mineral from the exporter to the importer in the year, zero if no trade is recorded	BACI	Thousand USD
POS	Equals 1 if TRADE is positive	BACI	Indicator
MFN	Rate applied by the importer to all partners together on the mineral in the year	WITS	%
PREF	Equals 1 if a PSR was retrieved for the exporter, importer, mineral and year	WITS	Indicator
RR	Realised rate, equal to the PSR where PREF is 1 and to MFN otherwise	WITS	%
DIFF	Equals 1 if RR differs from MFN	WITS	Indicator
RBT	Log of one plus RR, with RR in %	WITS	Logarithm
MFN-G	Rate of the importer carried in the gravity panel, equal for every exporter in an importer-mineral-year cell	WITS	%
HHI	Herfindahl-Hirschman index of destination concentration, the sum over the importers of the panel of the squared share of each importer in the exports of the exporter of the mineral in the year, zero if the exporter has no exports	BACI	Index, 0 to 1
Panel B: Gravity covariates			
AGR	Equals 1 if a regional trade agreement is in force between the pair, the value of the last observed year is carried forward after 2021	CEPII Gravity	Indicator
DIST	Log of the population-weighted bilateral distance in km	CEPII Gravity	Logarithm of km
CONTIG	Equals 1 if the pair shares a border	CEPII Gravity	Indicator
LANG	Equals 1 if the pair shares an official language	CEPII Gravity	Indicator
COL	Equals 1 if the pair shares a coloniser	CEPII Gravity	Indicator

Notes: BACI is the Base pour l'Analyse du Commerce International, WITS is the World Integrated Trade Solution and CEPII Gravity is the gravity database of the Centre d'Études Prospectives et d'Informations Internationales, rates from WITS are averaged with equal weights across the six-digit lines of a mineral listed in Table D5 and, for lithium, line 253090 (see the notes of Table D5), and partner groups that are not countries are not used as partner rates, PSR is the partner-specific rate, MFN-G and MFN are two records of the rate applied to all partners, they are equal in 55.8% of the rows of sample E and have a correlation of 0.93, only MFN is used to build the RBT, MFN-G enters the sample rule of sample E and the first panel of Table D4.

Table D3. Descriptive statistics.

Variable	Rows	Mean	Standard deviation	Minimum	Maximum
Panel A: Sample E					
TRADE	159,157	13,768.2	219,220.4	0	26,827,366
POS	159,157	0.303	0.460	0	1
RBT	159,157	1.000	0.768	0.000	3.761
RR	159,157	2.687	3.727	0.000	42.000
MFN	159,157	2.731	3.721	0.000	42.000
MFN-G	159,157	2.922	3.796	0.000	42.000
PREF	159,157	0.022	0.147	0	1
DIFF	159,157	0.016	0.127	0	1
AGR	159,157	0.284	0.451	0	1
DIST	159,157	8.596	0.886	5.159	9.894
CONTIG	159,157	0.045	0.207	0	1
LANG	159,157	0.062	0.242	0	1
COL	159,157	0.013	0.112	0	1
HHI	159,157	0.356	0.338	0.000	1.000
Panel B: Sample A					
MFN-G	186,383	2.954	3.776	0.000	42.000
AGR	186,383	0.282	0.450	0	1
DIST	186,383	8.584	0.891	5.159	9.894
HHI	186,383	0.353	0.338	0.000	1.000

Notes: Panel A reports the rows of sample E and Panel B the rows of sample A, variables are defined in Table D2 and samples in Table D1, TRADE is in thousand USD, RR, MFN and MFN-G are in %, RBT and DIST are logarithms, PREF is the indicator that a PSR was retrieved, the minimum and maximum of indicators are 0 and 1.

Table D4. Residual variance share of the policy regressors.

Variable	Varies along	(1)	(2)	(3)	(4)
Panel A: Sample A					
MFN-G	importer-mineral-year	0.2441	0.2441	0.0000	0.0000
AGR	pair-year	0.4940	0.1267	0.4940	0.1266
HHI	exporter-mineral-year	0.7318	0.7318	0.0000	0.0000
DIST	pair, time-invariant	0.5281	0.0000	0.5280	0.0000
Panel B: Sample E					
RBT	exporter-importer-mineral-year	0.4135	0.4037	0.0346	0.0227
MFN	importer-mineral-year	0.1743	0.1743	0.0000	0.0000

Notes: Entries are the RVS, the share of regressor variance that remains after the FE of each structure are removed, values below 1e-10 are shown as 0.0000, structure (1) is EY and IY effects, (2) adds PAIR effects, (3) is EMT and IMT effects and (4) adds PM effects, EY, IY, EMT and IMT are exporter-year, importer-year, exporter-mineral-year and importer-mineral-year effects, PAIR is an exporter-importer effect common to all minerals and PM is an exporter-importer effect specific to each mineral, Panel A uses the 186,383 rows of sample A and Panel B the 159,157 rows of sample E, variables are defined in Table D2 and samples in Table D1.

Table D5. Product lines and composition of sample E by mineral.

Mineral	Stages	HS 1992 codes	Rows	Share of zero flows	Mean RR (%)	Share with DIFF
Cobalt	Raw, processed and battery	260500, 282200, 810520, 285390	31,860	0.7924	2.15	0.0166
Copper	Raw and processed	260300, 740311	31,860	0.5796	0.99	0.0131
Graphite	Raw, processed and battery	250410, 380110, 854519	31,860	0.5670	3.14	0.0212
Lithium	Raw and processed	282520, 283691	31,812	0.7929	3.67	0.0167
Rare earths	Processed only	280530, 284690	31,765	0.7517	3.49	0.0145
All minerals	Five minerals pooled	13 lines in all	159,157	0.6967	2.69	0.0164

Notes: Rows are the rows of sample E, the lines are those of the six-digit 1992 nomenclature of the HS attributed to each mineral in the trade data, the codes were checked against the product descriptions of the BACI database when the list was built and 7 further candidate codes were removed, 2 because they were mapping errors and 5 because they were residual categories too broad to attribute to one mineral, two lines for batteries, 850760 and 850790, were retrieved but are not attributed to any one mineral and 253090 covers lithium-bearing and rare-earth-bearing substances that cannot be separated at six digits, so none of the three is in the trade data of the five minerals, the lithium tariff nevertheless averages the rates of 253090 and of the two lines listed, and without 253090 the coefficient of the RBT in column (8) of Table 1 is −0.3185 instead of −0.3257, RR is the realised rate and DIFF the indicator that it differs from MFN (Table D2).

Table D6. Importers and exporters in sample E.

Importer	Rows	PSR rows	DIFF rows	Importer	Rows	PSR rows	DIFF rows
Panel A: Importing countries							
United Arab Emirates	5,370	40	40	Greece	5,850	0	0
Austria	5,730	0	0	India	6,670	190	190
Belgium	5,730	0	0	Italy	5,850	0	0
Bulgaria	6,340	13	11	Japan	6,905	379	247
Brazil	6,905	456	456	Korea, Republic of	6,905	280	279
Canada	6,905	168	143	Malaysia	6,670	0	0
Switzerland	7,050	717	0	Netherlands	5,730	0	0
China	6,670	271	271	Poland	6,905	40	40
Germany	5,730	0	0	Sweden	5,730	0	0
Spain	5,730	0	0	Thailand	6,085	30	30
Finland	5,730	0	0	Türkiye	6,425	244	244
France	5,730	0	0	United States	5,715	530	530
United Kingdom	4,540	42	42	Viet Nam	5,557	97	92
Panel B: Exporting countries							
ARG, ARM, AUS, AUT, BEL, BGR, BRA, CAN, CHL, CHN, COD, COG, DEU, ECU, ESP, FIN, FRA, GBR, GEO, IDN, IND, IRN, JPN, KAZ, KOR, LAO, MEX, MMR, MNG, MRT, MYS, NLD, NOR, PAN, PER, PHL, PNG, POL, PRT, RUS, SAU, SGP, SRB, SWE, TUR, USA, UZB, ZAF, ZMB							

Notes: Importers are the 26 importers of sample E, rows are the rows of sample E, PSR rows are rows with a PSR retrieved and DIFF rows are rows whose RR differs from MFN, the 3,497 rows with a PSR are 2.2% of the 159,157 rows of sample E and 2,615 of them carry a rate that differs from MFN, exporters are the 49 exporters of sample E and are listed by their ISO 3166-1 alpha-3 codes.

Table D7. Rows with a partner-specific rate by year.

Year	PSR rows	Year	PSR rows	Year	PSR rows	Year	PSR rows	Year	PSR rows
1995	20	2001	27	2007	112	2013	76	2019	268
1996	10	2002	37	2008	132	2014	218	2020	328
1997	21	2003	48	2009	151	2015	217	2021	394
1998	11	2004	46	2010	124	2016	218	2022	0
1999	40	2005	44	2011	98	2017	278	2023	0
2000	39	2006	103	2012	145	2018	292	-	-

Notes: Entries are the rows of sample E with a PSR retrieved, the total is 3,497 rows, no PSR was retrieved for any row in 2022 or 2023, a hyphen marks an empty cell of the layout.

Appendix E. Functional form and reliability results

Table E1 collects the numerical results of the functional-form analysis and of the reliability checks discussed in the main text.

Table E1. Functional form, effect structures and reliability checks.

Panel A: Functional form

Regressor	Coefficient	Cut from 10% to 5% (%)	Cut from 5% to 0% (%)	Elasticity at mean rate	Observations
RBT (Table 1)	-0.3257*** (0.1147)	22 [6, 40]	79 [20, 168]	-0.24	89,994
RBT with the rate as a fraction	-16.3358*** (6.2190)	114 [21, 277]	122 [22, 302]	-0.43	89,994
Rate in levels (%)	-0.1601*** (0.0613)	123 [22, 306]	123 [22, 306]	-0.43	89,994

Panel B: Monte Carlo

Share of cells with a PSR (%)	RVS (%)	Mean estimate	RMSE	Mean SE	Coverage (%)
1	1.0	-0.520	0.194	0.157	87.2
2.2	2.2	-0.505	0.121	0.098	89.3
5	4.5	-0.513	0.072	0.063	91.0
10	7.6	-0.510	0.053	0.047	92.4

Panel C: Restricted samples

Sample	Coefficient	z	p-value	Observations	Rows before drops
Full sample (Table 1, column 8)	-0.3257*** (0.1147)	-2.84	0.005	89,994	159,157
Cells with a partner record	-0.2318 (0.1546)	-1.50	0.134	16,549	36,062
Importers with a partner rate	-0.2265* (0.1302)	-1.74	0.082	51,914	94,947

Panel D: Pair bootstrap

Inference	Estimate	SE	Lower limit	Upper limit	Share of draws below zero
Pair-clustered	-0.3257	0.1147	-0.5505	-0.1008	-
Pair bootstrap	-0.3403	0.1705	-0.6449	0.0113	0.970

Panel E: Effective clusters

Clustering	Clusters	Effective clusters (G*)	SE	95% interval, normal	95% interval, t on G* minus one
Directed pair	1,142	60.76	0.1147	[-0.5505, -0.1008]	[-0.5552, -0.0962]
Importer	26	7.40	0.1121	[-0.5454, -0.1059]	[-0.5960, -0.0554]
Exporter	49	9.83	0.1179	[-0.5568, -0.0945]	[-0.5932, -0.0581]
Mineral	5	1.16	-	-	-

Panel F: Effect structures

Structure	RBT	RVS of the RBT	MFN	RVS of MFN	Observations
EMT, IMT and PM (Table 1, column 8)	-0.3257*** (0.1147)	0.0227	-	0.0000	89,994
EMT, IMT and PAIR	-0.2813** (0.1144)	0.0248	-	0.0000	107,333
EY, IY, MY and PM	-0.2236** (0.0941)	0.0530	-0.1681 (0.1840)	0.0320	96,736
EY, IY, MY and PAIR (Table 1, column 7)	-0.2834 (0.1803)	0.1938	-0.6443** (0.2654)	0.1789	142,358
EY, IY and PM	0.0204 (0.1917)	0.0598	-0.3696 (0.3083)	0.0390	96,736

Notes: Significance levels: *** p<0.01, ** p<0.05, * p<0.1, SE in parentheses below each coefficient, SEs are pair-clustered, Panels A, C, D and E use the specification of column (8) of Table 1, RBT is the log of one plus the realised rate in % except in the second row of Panel A where the rate enters as a fraction, in Panel A the change in trade is exp(coefficient × change in the regressor) − 1 with the interval implied by the 95% interval of the coefficient in brackets and the elasticity is evaluated at the mean realised rate of 2.7%, Panel B has 200 replications per row with a true elasticity of −0.5, coverage is the share of replications whose 95% interval based on pair-clustered SEs contains it, RMSE is the root mean squared error, estimates in Panel B condition on identification, the importers sample in Panel C uses a solver tolerance of 1e-4, the bootstrap in Panel D resamples directed pairs and summarises 101 of 106 refits that converged, a hyphen marks an entry that is absorbed or not defined, which in Panel E reflects that five clusters are too few for a reliable cluster-robust variance, in Panel F MY is a mineral-year effect, PAIR is an exporter-importer effect common to all minerals and MFN replaces the RBT in a separate regression.

References

Ali, S. H., Giurco, D., Arndt, N., Nickless, E., Brown, G., Demetriades, A., Durrheim, R., Enriquez, M. A., Kinnaird, J., Littleboy, A., Meinert, L. D., Oberhänsli, R., Salem, J., Schodde, R., Schneider, G., Vidal, O., & Yakovleva, N. (2017). Mineral supply for sustainable development requires resource governance. *Nature*, *543*(7645), 367–372. https://doi.org/10.1038/nature21359

Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *American Economic Review*, *69*(1), 106–116. https://www.jstor.org/stable/1802501

Anderson, J. E., & van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, *93*(1), 170–192. https://doi.org/10.1257/000282803321455214

Armington, P. S. (1969). A theory of demand for products distinguished by place of production. *IMF Staff Papers*, *16*(1), 159–178. https://doi.org/10.2307/3866403

Baier, S. L., & Bergstrand, J. H. (2007). Do free trade agreements actually increase members' international trade? *Journal of International Economics*, *71*(1), 72–95. https://doi.org/10.1016/j.jinteco.2006.02.005

Balázsi, L., Mátyás, L., & Wansbeek, T. (2018). The estimation of multidimensional fixed effects panel data models. *Econometric Reviews*, *37*(3), 212–227. https://doi.org/10.1080/07474938.2015.1032164

Baldwin, R., & Taglioni, D. (2007). Trade effects of the euro: A comparison of estimators. *Journal of Economic Integration*, *22*(4), 780–818. https://doi.org/10.11130/jei.2007.22.4.780

Baltagi, B. H., Egger, P., & Pfaffermayr, M. (2003). A generalized design for bilateral trade flow models. *Economics Letters*, *80*(3), 391–397. https://doi.org/10.1016/S0165-1765(03)00115-0

Belloni, A., Chernozhukov, V., & Hansen, C. (2014a). High-dimensional methods and inference on structural and treatment effects. *Journal of Economic Perspectives*, *28*(2), 29–50. https://doi.org/10.1257/jep.28.2.29

Belloni, A., Chernozhukov, V., & Hansen, C. (2014b). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, *81*(2), 608–650. https://doi.org/10.1093/restud/rdt044

Bergstrand, J. H. (1985). The gravity equation in international trade: Some microeconomic foundations and empirical evidence. *Review of Economics and Statistics*, *67*(3), 474–481. https://doi.org/10.2307/1925976

Boehm, C. E., Levchenko, A. A., & Pandalai-Nayar, N. (2023). The long and short (run) of trade elasticities. *American Economic Review*, *113*(4), 861–905. https://doi.org/10.1257/aer.20210225

Chaney, T. (2008). Distorted gravity: The intensive and extensive margins of international trade. *American Economic Review*, *98*(4), 1707–1721. https://doi.org/10.1257/aer.98.4.1707

Chiang, H. D., Rodrigue, J., & Sasaki, Y. (2023). Post-selection inference in three-dimensional panel data. *Econometric Theory*, *39*(3), 623–658. https://doi.org/10.1017/S0266466622000081

Congressional Research Service. (2026, April 23). *Section 232 national security tariffs on copper imports* (CRS Insight IN12614). https://www.congress.gov/crs-product/IN12614

Conte, M., Cotterlaz, P., & Mayer, T. (2022). *The CEPII gravity database* (CEPII Working Paper No. 2022-05). Centre d'Etudes Prospectives et d'Informations Internationales. https://www.cepii.fr/CEPII/en/publications/wp/abstract.asp?NoDoc=13432

Correia, S., Guimaraes, P., & Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *Stata Journal*, *20*(1), 95–115. https://doi.org/10.1177/1536867X20909691

Davis, P. (2002). Estimating multi-way error components models with unbalanced data structures. *Journal of Econometrics*, *106*(1), 67–95. https://doi.org/10.1016/S0304-4076(01)00087-2

Eaton, J., & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, *70*(5), 1741–1779. https://doi.org/10.1111/1468-0262.00352

Egger, P., & Pfaffermayr, M. (2003). The proper panel econometric specification of the gravity equation: A three-way model with bilateral interaction effects. *Empirical Economics*, *28*(3), 571–580. https://doi.org/10.1007/s001810200146

Esteve-Pérez, S., Gil-Pareja, S., & Llorca-Vivero, R. (2020). Does the GATT/WTO promote trade? After all, Rose was right. *Review of World Economics*, *156*, 377–405. https://doi.org/10.1007/s10290-019-00367-w

Fally, T. (2015). Structural gravity and fixed effects. *Journal of International Economics*, *97*(1), 76–85. https://doi.org/10.1016/j.jinteco.2015.05.005

Farag, M., & Zaki, C. (2025). Allies and enemies: On the political determinants of trade in critical minerals. *Resources Policy*, *109*, 105701. https://doi.org/10.1016/j.resourpol.2025.105701

Feenstra, R. C. (2004). *Advanced international trade: Theory and evidence*. Princeton University Press.

Fernández-Val, I., & Weidner, M. (2016). Individual and time effects in nonlinear panel models with large N, T. *Journal of Econometrics*, *192*(1), 291–312. https://doi.org/10.1016/j.jeconom.2015.12.014

Fontagné, L., Guimbard, H., & Orefice, G. (2022). Tariff-based product-level trade elasticities. *Journal of International Economics*, *137*, 103593. https://doi.org/10.1016/j.jinteco.2022.103593

Freeman, R., Larch, M., Theodorakopoulos, A., & Yotov, Y. V. (2025). Unlocking new methods to estimate country-specific effects and trade elasticities with the structural gravity model. *Journal of Applied Econometrics*, *40*(6), 669–684. https://doi.org/10.1002/jae.3133

Gaulier, G., & Zignago, S. (2010). *BACI: International trade database at the product-level, the 1994-2007 version* (CEPII Working Paper No. 2010-23). Centre d'Etudes Prospectives et d'Informations Internationales. https://ideas.repec.org/p/cii/cepidt/2010-23.html

Gulley, A. L., Nassar, N. T., & Xun, S. (2018). China, the United States, and competition for resources that enable emerging technologies. *Proceedings of the National Academy of Sciences*, *115*(16), 4111–4115. https://doi.org/10.1073/pnas.1717152115

Head, K., & Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, & K. Rogoff (Eds.), *Handbook of international economics* (Vol. 4, pp. 131–195). Elsevier. https://doi.org/10.1016/B978-0-444-54314-1.00003-3

Head, K., & Ries, J. (2001). Increasing returns versus national product differentiation as an explanation for the pattern of US-Canada trade. *American Economic Review*, *91*(4), 858–876. https://doi.org/10.1257/aer.91.4.858

Helpman, E., Melitz, M., & Rubinstein, Y. (2008). Estimating trade flows: Trading partners and trading volumes. *Quarterly Journal of Economics*, *123*(2), 441–487. https://doi.org/10.1162/qjec.2008.123.2.441

Henderson, D. J., Soberón, A., & Rodriguez-Poo, J. M. (2022). Nonparametric multidimensional fixed effects panel data models. *Econometric Reviews*, *41*(3), 321–358. https://doi.org/10.1080/07474938.2021.1957283

International Energy Agency. (2026). *Global critical minerals outlook 2026*. IEA. https://www.iea.org/reports/global-critical-minerals-outlook-2026

Jin, S., Lu, X., & Su, L. (2025). Three-dimensional heterogeneous panel data models with multi-level interactive fixed effects. *Journal of Econometrics*, *249*, Article 105957. https://doi.org/10.1016/j.jeconom.2025.105957

Jochmans, K. (2017). Two-way models for gravity. *Review of Economics and Statistics*, *99*(3), 478–485. https://doi.org/10.1162/REST_a_00620

Kalantzakos, S. (2020). The race for critical minerals in an era of geopolitical realignments. *International Spectator*, *55*(3), 1–16. https://doi.org/10.1080/03932729.2020.1786926

Kapetanios, G., Serlenga, L., & Shin, Y. (2021). Estimation and inference for multi-dimensional heterogeneous panel datasets with hierarchical multi-factor error structure. *Journal of Econometrics*, *220*(2), 504–531. https://doi.org/10.1016/j.jeconom.2020.04.011

Korinek, J., & Kim, J. (2011). Export restrictions on strategic raw materials and their impact on trade and global supply. *Journal of World Trade*, *45*(2), 255–281. https://doi.org/10.54648/trad2011009

Larch, M., Shikher, S., Syropoulos, C., & Yotov, Y. V. (2021). *Quantifying the impact of economic sanctions on international trade in the energy and mining sectors* (CESifo Working Paper No. 8878). CESifo. https://hdl.handle.net/10419/232475

Larch, M., Wanner, J., Yotov, Y. V., & Zylkin, T. (2019). Currency unions and trade: A PPML re-assessment with high-dimensional fixed effects. *Oxford Bulletin of Economics and Statistics*, *81*(3), 487–510. https://doi.org/10.1111/obes.12283

Lu, X., & Su, L. (2026). Mundlak estimators for three-dimensional panel data models. *Economics Letters*, *262*, Article 112842. https://doi.org/10.1016/j.econlet.2026.112842

MacKinnon, J. G., Nielsen, M. O., & Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, *232*(2), 272–299. https://doi.org/10.1016/j.jeconom.2022.04.001

Mayer, T., & Zignago, S. (2011). *Notes on CEPII's distances measures: The GeoDist database* (CEPII Working Paper No. 2011-25). Centre d'Etudes Prospectives et d'Informations Internationales. https://doi.org/10.2139/ssrn.1994531

Pfaffermayr, M. (2026). Bias-corrected cluster-robust standard errors for fixed effects PPML estimators of gravity panel models with autocorrelated disturbances. *Empirical Economics*, *70*(2), Article 31. https://doi.org/10.1007/s00181-026-02888-4

Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, *88*(4), 641–658. https://doi.org/10.1162/rest.88.4.641

Santos Silva, J. M. C., & Tenreyro, S. (2011). Further simulation evidence on the performance of the Poisson pseudo-maximum likelihood estimator. *Economics Letters*, *112*(2), 220–222. https://doi.org/10.1016/j.econlet.2011.05.008

Santos Silva, J. M. C., & Tenreyro, S. (2022). The log of gravity at 15. *Portuguese Economic Journal*, *21*(3), 423–437. https://doi.org/10.1007/s10258-021-00203-w

Sasaki, Y., & Ura, T. (2026). Slow movers in panel data. *Econometric Theory*. Advance online publication. https://doi.org/10.1017/S0266466625100157

Snoussi-Mimouni, M., & Avérous, S. (2024). High demand for energy-related critical minerals creates supply chain pressures. *WTO Data Blog*. https://www.wto.org/english/blogs_e/data_blog_e/blog_dta_10jan24_e.htm

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, *58*(1), 267–288. https://doi.org/10.1111/j.2517-6161.1996.tb02080.x

United Nations Conference on Trade and Development. (2026a, June 12). *Critical minerals are reshaping global trade as demand surges*. UNCTAD. https://unctad.org/news/critical-minerals-are-reshaping-global-trade-demand-surges

United Nations Conference on Trade and Development. (2026b). *Global trade update (June 2026): The shifting dynamics of critical minerals trade* (UNCTAD/DITC/INF/2026/6). UNCTAD. https://unctad.org/publication/global-trade-update-june-2026-shifting-dynamics-critical-minerals-trade

Vivoda, V. (2023). Friend-shoring and critical minerals: Exploring the role of the Minerals Security Partnership. *Energy Research & Social Science*, *100*, Article 103085. https://doi.org/10.1016/j.erss.2023.103085

Watari, T., Nansai, K., & Nakajima, K. (2020). Review of critical metal dynamics to 2050 for 48 elements. *Resources, Conservation and Recycling*, *155*, Article 104669. https://doi.org/10.1016/j.resconrec.2019.104669

Weidner, M., & Zylkin, T. (2021). Bias and consistency in three-way gravity models. *Journal of International Economics*, *132*, Article 103513. https://doi.org/10.1016/j.jinteco.2021.103513

Yang, Y., & Zhang, H. (2023). Three-way gravity models with multiplicative unobserved effects. *The Econometrics Journal*, *26*(3), 422–443. https://doi.org/10.1093/ectj/utad012

Yotov, Y. V., Piermartini, R., Monteiro, J.-A., & Larch, M. (2016). *An advanced guide to trade policy analysis: The structural gravity model*. UNCTAD and WTO. https://unctad.org/publication/advanced-guide-trade-policy-analysis-structural-gravity-model-volume-2

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- manuscriptmineralselasticitysupplementv33.docx