

Policy transmission and institutional resilience

Exposure, outside options and the incidence of restrictions

Elies Lucio

Calibration 2 · 4 October 2026

Abstract

This global edition separates recorded trading routes, effective regulatory exposure and institutional response for gold, cobalt, tantalum, tin and tungsten. It preserves the original African diagnostic and extends the evidence archive to worldwide origin and destination margins for 2005–2024, with selected bilateral 3T routes. The established household and gate mechanisms produce 1,120 allocations across zero and concave external retention scenarios, reproducing the original zero-retention benchmark exactly. Normalized village experiments are not scaled into country estimates without buyer and household data. A separate nonlinear reconstruction implements genuine VCG seller retention and Muir resale, allowing shared conclusions and mechanism-dependent incidence to be distinguished. Country classifications remain descriptive and unmeasured parameters remain explicit scenarios; the extension supplies no causal significance claims.

Status. Global extension of the preserved conditional calibration. Numerical and accounting verification does not establish a causal policy effect or administrative approval. Monetary and temporal units are defined in the text.

Scope of this edition

This edition adds the global comparison and country appendices to the preserved benchmark argument. The original gold/cobalt sections retain their original experimental populations and reported numbers; they are reproduced benchmarks, not results of the enlarged country grid. The sections headed “Global” or “Five-mineral” report the extension. This distinction is essential for cobalt, whose newly explicit world producers change the assumed candidate composition. The three alternative nonlinear papers are delivered separately in **consolidated calibration**.

1 Institutions determine where a restriction is felt

A buyer restriction reaches extraction through the institutions connecting a sale to a miner’s livelihood. Those institutions determine who can find another purchaser, who finances necessities, and which households can leave mining. A restriction may lower an intermediary’s profit while leaving output unchanged. It may eliminate a certified sale while most production continues outside certification. It may also improve a local institutional outcome without producing a measurable decline in national mineral tonnage. These possibilities concern different margins, and their empirical evaluation requires different denominators.

The quantitative distinction is substantial. In the screening exercise examined here, the share of displaced certified entries accounted for by lower extraction can approach 90 percent, yet aggregate candidate extraction falls by only 0.016 percent. The reason is the small quantity

initially displaced. In the separate village exercise, all three institutional mechanisms together absorb the entire contraction of a restricted benchmark, but that benchmark contraction is itself small and the village primitives are hypothetical. Neither result establishes the effectiveness or ineffectiveness of an existing regulatory framework. They establish mechanisms and magnitudes conditional on their stated environments.

This contribution separates three layers of evidence. Archived trade records describe changes in the destination composition of recorded imports. A common-market screening calibration supplies a conditional connection between lost certified entries, outside sales and extraction. A normalized village model compares selling institutions, an integrated goods-and-mineral interface and subsistence dependence. The layers are informative together because they expose missing links in the transmission chain. They are not independent empirical confirmations of one policy effect. Their shared concern is how a restriction changes the marginal reward to production, after allowing the relevant institutions to respond.

2 From destination shares to effective exposure

Trade geography is an observable starting point. It is not effective policy reach. A shipment’s recorded destination need not be its final consumer, its first demanding buyer or the jurisdiction whose rule changes its acquisition cost. Refining, processing and re-export can separate these stages. Enforcement can vary among firms within a destination, and a nominally uncovered destination can contain buyers implementing private sourcing standards. Consequently, a destination share is neither automatically a point estimate nor a general upper or lower bound for the expenditure exposed to a restriction.

The archived cobalt comparison illustrates both the value and the limit of geographic evidence. For HS810520, covering cobalt intermediates, unwrought metal and powders, the DRC’s recorded EU share rises from 0.0105 percent in the pooled 2015–2017 baseline to 0.1345 percent in 2024. China’s corresponding share is 99.3639 and 98.9814 percent. For all African origins, the EU share declines from 2.4808 to 0.8599 percent, the US share from 4.9763 to 1.6985 percent and China’s share rises from 85.2573 to 92.6738 percent. These numbers divide recorded partner-import value by the same product’s total returned destination value. They do not measure fine cobalt mass, illegal extraction or final regulatory coverage. ([United Nations Statistics Division, 2026](#))

Table 1: Recorded cobalt destination shares

HS810520 origin scope	Destination	Pooled 2015–2017 share (%)	Recorded 2024 share (%)
DRC	EU27	0.0105	0.1345
DRC	China	99.3639	98.9814
African origins	EU27	2.4808	0.8599
African origins	United States	4.9763	1.6985
African origins	Canada	0.3623	0.0504
African origins	China	85.2573	92.6738

Notes: Source: archived UN Comtrade responses, reconstructed from the archived responses in the replication tables. Available returned reporters; shares are nominal recorded value. A displayed zero elsewhere in the full table means no positive returned record, not verified absence of trade.

Even this descriptive African aggregate combines two changes: individual origins alter their destination shares, and the origins’ weights in recorded African exports change. An exact

symmetric decomposition for six origins with comparable positive denominators separates them. Within-origin routing changes raise the combined EU27/US/Canada/Australia share by 2.0876 percentage points; changing origin weights lower it by 7.2400 points. The cohort's net change is therefore -5.1524 points. The full available-origin aggregate changes by -5.2111 points, with a separate -0.0587-point bridge for differences between the full and comparable-origin populations. Reading the aggregate decline as every origin moving away from these destinations would miss this composition effect.




Figure 1: Origin composition reverses the interpretation of the aggregate change

Notes: HS810520; combined EU27, United States, Canada and Australia recorded import-value share. Six comparable African origins, pooled 2015–2017 versus 2024. The full-origin aggregate adds a separate -0.0587 percentage-point bridge. This is a descriptive identity, not a causal policy decomposition.

Historical projections require similar discipline. The cobalt headline comparison carries a three-year baseline seven years beyond its final observation; the other mineral table uses an eight-year horizon. One-year forecast errors would not assess that extrapolation. Matched-horizon descriptive checks on the same archived data produce pooled-share absolute errors averaging 26.14 percentage points for the DRC cobalt western-destination share and 24.44 points for African cobalt. The analogous African unwrought-gold figure is 7.71 points. A linear-share projection, constrained to valid shares, changes these errors substantially and performs better for some products and worse for others. These are overlapping historical windows already available during analysis, not fresh holdouts or prediction intervals. They establish the importance of benchmark choice without providing a causal counterfactual for the 2019 sourcing milestone or the 2021 EU application date.

3 Screening incidence and the scale of extraction

Let candidate extraction be Q , outside sales O and candidate entries into certification E . The physical account is $Q = O + E$. For a screening change from state 0 to state 1, define conditional incidence as $T = (Q_0 - Q_1)/(E_0 - E_1)$. Its denominator is the flow displaced from certification. Its complement measures replacement through increased outside sales when this two-channel account applies. Neither denominator is total national production, all certified sales or a no-institution benchmark.

In the maintained screening environment, outside expenditure pO is fixed and candidate supply follows $Q = Bp^\eta$. Writing $\delta = 1 - Q_1/Q_0$, the exact finite incidence is

$$T = \frac{\delta}{\delta + (O_0/Q_0)[(1 - \delta)^{-1/\eta} - 1]}. \quad (1)$$

The corresponding local expression is $\eta Q/(\eta Q + O)$. Most candidate output in the gold calibration already serves the outside market. The local incidence consequently resembles $\eta/(1 + \eta)$: the supply curvature strongly shapes how the loss of a small certified flow divides between reduced production and additional outside sales. This mechanism does not require treating outside-market production as verified illegal output.

All 22 screening cases compare assumed discovery increasing from .25 to .50 at unchanged joint inputs within the pair. The reference gold calculation reduces candidate output by 0.005913 percent, while 33.3424 percent of displaced certified entries becomes reduced extraction. The cobalt reference reduces candidate output by 0.286023 percent and has 33.8305 percent incidence. Different curvature and initial-quote assumptions change both quantities. Across the stated gold cases, extraction changes range from -0.016003 to -0.003327 percent; cobalt's range is -0.507554 to -0.156380 percent. These are explored conditional sensitivities, not confidence intervals or empirically identified policy-effect bounds.

Table 2: A high incidence share can apply to a small displaced flow

Screening case	Candidate extraction change (%)	Lost certified entries / initial candidate output (%)	Conditional incidence (%)
Gold reference	-0.005913	0.017735	33.3424
Gold, Colombia procurement transport	-0.012437	0.017765	70.0085
Gold, Peru lower procurement endpoint	-0.015111	0.017777	85.0052
Gold, Peru upper procurement endpoint	-0.016003	0.017781	90.0037
Cobalt reference	-0.286023	0.845459	33.8305

Notes: Policy: assumed discovery .25→.50. Gold is the common gold market, including Colombia and Peru; cobalt is separate. Country-named rows transport a regional procurement moment to the common gold curvature under zero net recovery. Statistical inference: NOT ESTIMABLE FROM CURRENT EVIDENCE.

The exact aggregate identity is $-(Q_1 - Q_0)/Q_0 = T(E_0 - E_1)/Q_0$. At the upper Peru-price endpoint, 90.0037 percent applies to a displaced flow of only 0.017781 percent of initial candidate production. Their product is the reported 0.016003 percent contraction. In the source scale, the explored gold changes are approximately -0.288 to -0.060 tonnes relative to 1,798.972 tonnes of baseline candidate output; cobalt changes are -307.578 to -94.767 tonnes relative to 60,600 tonnes. Those are common-market model strata, including assumed rest-world assignments. They cannot be presented as estimated national reductions in illegal mining. The source's 2024-reference physical vector is preserved as a named comparison rather than silently mixed with another vintage.

Procurement evidence provides a conditional connection to curvature. Write the procurement/world quote ratio as ζ and net implementation recovery relative to sales as m . In the implemented power-cost account, $\zeta = \eta/(1 + \eta) + m$. Setting $m = 0$ turns a .70 quote ratio into $\eta = 7/3$, .85 into $17/3$ and .90 into 9. This explains why the transported incidence values closely resemble the quote ratios; their numerical proximity is partly imposed by the mapping. Colombia's regional price-level evidence does not estimate a policy response, and Peru's .85-.90 interval

remains provisional with unresolved sampling, date and source-access limitations. (UNODC and MME, 2025; BCRP, 2026)

Allowing nonnegative recovery changes the conclusion. The admissible restriction is $0 < \zeta - m < 1$, and $\eta = (\zeta - m)/(1 - \zeta + m)$. If $0 \leq m < \zeta$, η can approach zero. At the fixed outside share, local incidence therefore has limiting lower bound zero and an upper value at zero recovery: approximately 70.01 percent for the Colombian ratio and 90.00 percent across the Peruvian interval. A positive lower bound requires further restrictions. If net recovery can be negative, those upper limits need not apply. These are conditional local mathematical statements, not new empirical intervals for either screening incidence or the paper’s VCG transmission coefficient.

4 Three institutions and eight distinct comparisons

The village exercise makes institutional adjustment explicit. A household allocates time between mining and leisure, buys an essential basket, and retains cash. Mining income, outside wealth and purchased necessities enter a joint budget. Participation requires at least the utility of a declared outside allocation. The intermediary chooses procurement terms, and, where the interface is available, a goods margin. Implementation cost reconciles procurement spending less goods-margin recovery with the resource cost of goods plus household cash less outside wealth, plus transport. Recovery follows the allocation; it is not imposed as an arbitrary proportional discount on extraction costs.

Three channels define eight configurations. Selling arrangements S permit the relevant Muir or VCG institution instead of fixed contracts. Interface I allows an integrated intermediary to choose a goods margin. Subsistence dependence H removes an independently delivered basic-consumption allowance. H therefore changes the household’s outside support and the procurement terms needed to induce supply; it is not a free welfare-neutral switch. Every configuration is compared with its own no-charge optimum. Forcing all eight to share the same baseline would mix the effect of the restriction with the effect of changing the institution.

Table 3: Each institution is compared with its own baseline

Configuration	Muir Q(0)	Muir Q(.1)	VCG Q(0)	VCG Q(.1)
Empty	0.129800	0.129715	0.169650	0.168850
S	0.130020	0.130020	0.162575	0.162575
I	0.137028	0.136940	0.176950	0.176350
H	0.318625	0.318625	0.318625	0.318625
S+I	0.137190	0.137190	0.170925	0.170925
S+H	0.318625	0.318625	0.318625	0.318625
I+H	0.318625	0.318625	0.324550	0.324550
S+I+H	0.318625	0.318625	0.324550	0.324550

Notes: Normalized mineral units in one hypothetical static period. All primitives are assumed. Values rounded for display; paired differences use full precision. The two selling institutions have different buyer economies, so their levels are not estimates of a common market’s institutional treatment effects.

For Muir, the benchmark extraction contraction is .000085, or .0655 units per 100 benchmark units. For VCG it is .000800, or .4716 per 100. Full-configuration output is unchanged in both examples. Absorption is defined as benchmark contraction minus full-configuration contraction, divided by benchmark contraction when that denominator is meaningful. The resulting 100 percent describes preservation of these small benchmark losses. It is not a claim that institutions neutralize every restriction or preserve every socially relevant outcome.

The mechanisms differ. On the supported Muir branch, the buyer charge subtracts the same amount from receipts for every feasible quantity above regulated buyer mass: its marginal effect on receipts with respect to quantity is zero. In the full base example, output stays fixed while profit falls by .00049. Fixed-route sales instead expose a constant fraction of every additional unit, lowering marginal receipts. Muir’s result is therefore a comparison of how the restriction enters the seller’s marginal problem, conditional on continued operation within the supported charge range.

VCG has a different source of insulation. The gate’s payment depends on the externality imposed on rival buyers. When a covered highest-value buyer remains ahead, a lower bid need not alter the losing bids determining that payment. Covering a losing buyer can change those externalities directly. Coverage of nominal sales thus does not translate mechanically into marginal gate-receipt exposure. Estimating that transmission would require the reserve, rival bids and relevant inventory margin, not just the covered buyer’s observed sales share.

Subsistence configurations add a participation boundary. The base solutions with H lie close to the household’s participation floor. Their unchanged output is supported by that boundary, not by a smooth interior theorem guaranteeing a zero response. Profit can absorb a small charge while the least costly participating allocation remains unchanged. A larger shock or a stronger outside option could instead change the active branch or make operation unprofitable. The reported examples cannot certify the behavior of those unexamined regions.

5 Joint sensitivities, channel rankings and economic importance

The expanded exercise declares five joint primitive bundles before computation, each evaluated at charges .05 and .10 in both institutions. The bundles combine outside consumption, wealth, subsistence requirements, goods-margin limits and buyer coverage rather than treating each uncertain quantity as an independent probability draw. Together they add 320 allocations to the 240 reproduced source allocations. Every declared case solves. The ranges are extrema over these finite experiments; intermediate primitive values and other closures are not certified by their endpoints.

In the added cases, the empty Muir benchmark changes by -.000175 to -.0000425 normalized units, or -.1271 to -.0319 percent of its own baseline. Empty VCG changes by -.001500 to -.000300 units, or -.8842 to -.1627 percent. Full-configuration changes remain zero in the explored set. A positive contraction target is therefore not attained by any of these full-configuration pairs. Without an empirical target or a calibrated mapping to a national population, selecting a universal acceptable contraction would add an unsupported policy judgment. The useful practical result is a frontier: any chosen target must be compared with the paired extraction changes, rather than with the absorption percentage or incidence ratio.

Attribution also depends on institutions. Shapley contributions average a channel’s incremental preservation of output over all six orders in which S, I and H can be introduced. They retain interactions and sum to total absorbed quantity. For base Muir, the selling and subsistence contributions each equal .03306 per 100 empty-baseline units, while the interface contribution is approximately -.00064. The small latter magnitude is resolution-sensitive. For base VCG, selling contributes .21613, the interface .03930 and subsistence .21613 per 100 empty-baseline units; finer physical packages reduce the interface term to .01965. These figures do not support a precise general ranking of the interface channel.

Covering a losing VCG buyer makes the selling contribution negative in both the original sensitivity and the added joint scenarios. Selling institutions can then increase the policy contraction relative to a preceding configuration, while other channels still preserve the full model's output. The sign is economically meaningful and is retained. Likewise, eliminating the permissible goods margin makes the interface redundant, and doubling Muir's fixed-route regulated share limits its capacity to .25. Some subsistence configurations then operate at zero output both before and after the charge. Their zero change records inactivity, not successful preservation of a functioning mine. Own-baseline percentages are left undefined where that baseline is zero.

Numerical accuracy matters most for small channel differences. The original finer grid changes selected Muir output by at most .00000125 and VCG output by .000025. The VCG benchmark contraction becomes .000750 rather than .000800. This is package-granularity sensitivity, not empirical uncertainty. Budget, physical inventory and consolidated cost identities hold at numerical precision across all 560 allocations. These checks support the conditional calculations while leaving their hypothetical economic premises visible.

6 What remains to be measured

No statistical confidence interval or p-value is estimable for these new structural policy effects from the current evidence. The village primitives are assumed; the common-market screening change is not an identified intervention; and trade totals do not reveal the relevant treatment assignment or counterfactual. The proper entry is **NOT ESTIMABLE FROM CURRENT EVIDENCE**, not "statistically insignificant." Counting model cases, transaction templates or overlapping historical forecasts as independent observations would manufacture an inferential design.

External studies can inform the institutional question without filling that gap automatically. The published eastern DRC programme study reports matched estimates for 3T sites, including army presence, taxation requests and cassiterite production. Its annual cassiterite-production estimate is -1.290 tonnes per mine with standard error 11.599 tonnes and reported $p=.913$; the mine sample contains 58 observations in 43 clusters. That result is not an equivalence finding, a cobalt estimate or a p-value for this village model. Different institutional outcomes can respond differently, and a national illegal-output conclusion cannot be recovered by importing an unrelated study's significance level. (PRG et al., 2020)

The prepared panel contains fifteen empty table templates, not observations. A defensible next empirical stage would link sampled sites, physical inventories, lots, assays, sellers, buyer bids, household accounts and dated framework exposure. The sampling frame must retain inactive and exiting sites. Mineral content, moisture, recovery, deductions, credit and payment delay must be measured at compatible transaction stages. Coverage must follow the buyer and obligation actually affecting the transaction; independent observations must establish whether the rule was applied and enforced.

Household measurement is equally central. Purchased and independently supplied necessities, outside income, wealth, debt, mining hours and feasible alternatives determine whether subsistence binds and whether a procurement quote supports participation. An observed markup cannot by itself identify net recovery after transport, credit losses and other costs. A price level cannot identify the supply response without variation and an implementation account. To estimate policy effects, a prospectively specified assignment or credible comparison design must preserve

common shocks and market interference, and carry shared measurement and fitting uncertainty through the paired calculations.

Rare earths remain a feasibility branch. Empty narrow refined-product queries do not demonstrate absent mining or absent concentrate trade. A calibration would first require a named mineral species and product stage, representative assays, recovery and payability, matched buyers and locally observed production and household accounts. Mixed ore headings cannot supply those objects. The present result is consequently a rigorous measurement agenda alongside completed conditional mechanisms: institutions can redirect the incidence of a restriction, but their empirical ability to preserve extraction depends on exposure, participation and market information that the current archive does not yet identify.

7 The five-mineral evidence base

Gold, cobalt, tantalum, tin and tungsten clear in five separate world markets. The reference year is 2024. Coltan denotes an ore or concentrate; the tantalum model uses contained Ta. Gross concentrates, oxide and contained-metal mass are not interchangeable. Prices remain attached to their stated stage: converting an oxide-equivalent quotation into dollars per tonne of metal does not turn it into a measured minegate receipt. This is particularly important for the 3T minerals. (USGS, 2026; BGS, 2026)

The common archive contains 6,277 BGS country–product–year records for 2005–2024 across the five core minerals and copper, nickel and diamonds. Its global country panel is a separately versioned source scenario, not an average of incompatible USGS and BGS observations. Notes identify export proxies, fiscal years and uncertain content. Gross or niobium-containing tantalum records with no valid Ta conversion remain unconverted. National scales and world residuals are reconciled in the structural registry; the BGS source panel serves the wider descriptive comparison. All observations retain their provenance.

UN Comtrade extraction covers each year from 2005 to 2024, with 13,125 destination margin records, 14,362 origin margin records and 1,021 records in a six-country 3T route panel. These are the declared query populations, not a complete global bilateral network. Missing and zero observations are distinct. The World partner is an aggregate, not an extra country or territory; origin and destination margins must not be added to each other. HS261590 combines niobium, tantalum and vanadium. Without finer national evidence it remains a mixed trade indicator, never pure tantalum extraction or a direct regulatory-exposure measure. HS2615 is broader still and also includes zirconium. (UN Comtrade, 2026; UN Statistics Division, 2017)

IPIS contributes dated, visited mine-site and institutional observations in DRC and CAR. The DRC raw file contains 7,163 records for 3,423 sites; the admitted historical descriptive subset uses onsite IPIS-source visits through 2024. Administrative qualifications and later visits are not silently pooled with that population. Selection into a visit prevents a simple visited-site fraction from identifying the national candidate share. EITI workbooks for 13 country contexts describe fiscal institutions; their often earlier vintages and aggregate currency accounts do not identify 2024 commodity-specific collection. (IPIS, 2026; EITI, 2026)

Rough diamonds provide a certification comparison through the Kimberley Process. Carats, values and certificate records are kept separate: quality and sorting prevent direct insertion into the homogeneous-metal model. Copper and nickel describe industrial and coproduct contexts. Their prices and output do not enter the five-mineral equilibrium unless an explicitly named

joint-revenue experiment provides the physical and revenue accounts. Observed 2025 updates remain separate from the 2024 calibration. (Kimberley Process, 2026)

Table 4: Reference scales and fixed-value weights, 2024

Mineral	Mine content (t)	Reference USD/t	Weight (%)
Gold	3,661.2	76,718,111	93.39
Cobalt	302,000.0	26,103	2.62
Tantalum	2,500.0	203,915	0.17
Tin	294,000.0	30,159	2.95
Tungsten	82,000.0	31,778	0.87

Notes: Weights are $P_{m,2024}Q_{m,2024}/\sum_j P_{j,2024}Q_{j,2024}$. Gold retains the original WGC scale; other core minerals use the admitted USGS vintage. Prices concern different product stages. The resulting index is an accounting comparison, not mine revenue or welfare.

8 Global routing and the institutional evidence boundary

The global archive extends the original African trade diagnostic while retaining 2007–2024 as its historical comparison window and adding 2005–2006 where available. Destination shares are descriptive trade-value shares. A declared destination does not identify the ultimate buyer, covered use, beneficial owner, customs enforcement or the effective reach of a restriction. Re-exports, mirror differences and mixed HS classifications remain part of the evidence record.

The 3T country priorities combine production, trade and institutional roles: DRC and Rwanda lead tantalum coverage, followed by Nigeria, Ethiopia and Mozambique; Indonesia and Bolivia join the tin cases, with explicit limitations for Myanmar; and Rwanda and Bolivia are named tungsten structural origins, with DRC retained in the BGS and IPIS descriptive evidence. Burundi and Uganda are examined separately as possible production and trading cases. Brazil, Australia, China and other industrial producers supply the world comparators. These lists do not classify countries or industries as illegal.

Where paired trade years are observed, the routing identity separates within-origin changes from composition. With origin export weights w_i and destination shares s_i , $S = \sum_i w_i s_i$ and

$$\Delta S = \sum_i \bar{w}_i \Delta s_i + \sum_i \bar{s}_i \Delta w_i, \quad \bar{z}_i = (z_{i,0} + z_{i,1})/2.$$

An origin absent from an otherwise incomplete source is not assigned a zero share. Global source margins and the narrower bilateral diagnostic have different coverage, which is reported with the denominator. Neither component is a causal effect of certification.

9 Institutional experiments with zero and concave retention

The extension computes 70 mechanism–scenario–retention contexts and 1,120 allocations. Each of the eight combinations of selling flexibility, goods interface and subsistence is reoptimized under the policy and its own baseline. All six channel orderings enter the Shapley decomposition. The original zero-retention benchmark reproduces 1,920 checked cells exactly.

The additional inventory value is $b(s) = 0.30s - 0.5s^2$ in normalized units. It is an external value to the gate from stock not offered. It is not the seller’s reserve within the legacy VCG

auction. The flexible gate optimizes the offered quantity with the preserved Muir receipt envelope or whole-buyer VCG receipt schedule; fixed contracts retain their complete-offer rule. Receipt cash and retained-stock value remain separate. There are 100 allocations with positive retained stock in the full design, with a maximum of approximately 0.16073 normalized units.

Table 5: Institutional baseline comparisons in normalized mineral units

Mechanism	Retention	Fixed-config. fall	Full-config. fall	Absorbed share
Muir	zero	0.0000850	0.0000000	1.000
Muir	concave	0.0000850	0.0000000	1.000
VCG	zero	0.0008000	0.0000000	1.000
VCG	concave	0.0008000	0.0000000	1.000

Notes: Base household and buyer scenario; the quantities are model units, not tonnes. Absorption is relative to each mechanism and retention case’s own fixed-configuration response.

These village exercises supply a measurement blueprint applicable to each core mineral. They are not multiplied by national tonnage to manufacture household estimates. Mine visits and export accounts do not identify the required joint distribution of household outside options, productivity, buyer capacity, goods recovery or bargaining terms. The Rwanda deduction records support the relevance of interlinked payments but do not estimate those primitives.


The nonlinear companion has a genuine seller-retention VCG mechanism and a separately optimized Muir resale economy. Its institutional experiments therefore test a different construction. Agreement in sign is classified as shared under related assumptions; changed channel rankings, payment incidence or inventories are mechanism-dependent. No numerical similarity is described as independent empirical confirmation.

Table 6: Recorded routing toward China, 2005–2024

Product	Change (pp)	Within (pp)	Mix (pp)	Retained value (%)
Tin	63.41	55.02	8.39	97.4 / 100.0
Tungsten	61.49	49.33	12.16	100.0 / 100.0
Mixed Nb/Ta/V	44.75	47.00	-2.25	85.5 / 74.3

Notes: Six requested origins: DRC, Rwanda, Nigeria, Indonesia, Bolivia and Myanmar. Only common reporters and origins with positive recorded endpoint values enter the identity. The last column gives retained recorded value in 2005 and 2024. Mixed Nb/Ta/V excludes DRC, Bolivia and Myanmar from the balanced comparison; tungsten excludes Indonesia. This is not a full global bilateral estimate.

Observed customs destination concentration




Source: UN Comtrade, annual recorded imports. Reporter coverage varies by year. Missing flows are not filled with zero. Each mineral uses its declared HS basket.

HS261590 mixes Nb/Ta/V. Destination concentration is not effective regulatory exposure.

Figure 2: Recorded global routes and selected-origin decomposition

Notes: Recorded trade values and common-support restrictions are explicit. Destination is not effective regulatory exposure.

Selected six-origin 3T routes to China



Origins queried: DRC, Rwanda, Nigeria, Indonesia, Bolivia and Myanmar.
Only common reporters and positive-origin support enter the endpoint decomposition.

The mixed Nb/Ta/V comparison excludes DRC, Bolivia and Myanmar after matching.
These nominal-value accounting changes do not identify physical or causal re-routing.

Figure 3: Recorded global routes and selected-origin decomposition

Notes: Recorded trade values and common-support restrictions are explicit. Destination is not effective regulatory exposure.

The destination shift is an exact nominal-value accounting decomposition. Within-origin changes can reflect price, grade, reporting and routing as well as behavior. Changing origin weights need not reflect new mining. The mixed Nb/Ta/V restriction materially changes the sample: its unrestricted recorded China shares are approximately 44.00% and 89.77%, whereas the common-support shares are 51.46% and 96.21%. No missing route is interpreted as zero actual trade.

A The hypothetical village economy

All quantities and currency amounts in this appendix are normalized for one static village period. The following primitives are assumptions, not estimates of a country or mineral. Let household labour be $\ell \in [0, 1)$, mineral output $q = A\ell$, purchased necessities $x \geq 0$, independently delivered necessities t , retained cash $b \geq 0$, and outside wealth a . At procurement price θ and goods margin μ , the household solves

$$\max_{\ell, x, b} U = \log(t + x - c) + \eta_h \log(1 - \ell) + \beta_h \log(b + b_0), \quad (2)$$

$$(b_z + \mu)x + b = a + \theta A\ell, \quad t + x > c. \quad (3)$$

Here c is basic consumption and b_z is the basket's resource price. Participation requires

$$U \geq \bar{U} = \log s_o + \eta_h \log L_o + \beta_h \log(b_o + b_0), \quad (4)$$

where (s_o, L_o, b_o) is a separately declared outside allocation: consumption above needs, leisure and cash. The independent allocation is available upon exit; it is not a transfer conditional on accepting the gate's contract.

Table 7: Assumed base primitives of the village experiment

Primitive	Meaning	Value
A	Mineral output per unit of labour	0.4
(c, a)	Basic consumption; outside wealth	(0.05, 0.02)
(η_h, β_h, b_0)	Leisure weight; cash weight; cash offset	(1, 0.25, 0.005)
b_z	Resource price of necessities	1
(s_o, L_o, b_o)	Outside surplus consumption, leisure and cash	(0.008, 0.25, 0.001)
$(\theta_{\min}, \theta_{\max}, \bar{\mu})$	Procurement bounds; maximum goods margin	(0.001, 2, 0.5)
(t_1, t_2)	Transport coefficients, $\mathcal{T}(q) = t_1 q + t_2 q^2/2$	(0.01, 0.25)
τ	Covered-buyer charge per mineral unit	0.10

The three switches are explicit. With interface I absent, $\mu = 0$; with I present, μ is chosen from 41 equally spaced quotes on $[0, \bar{\mu}]$. With subsistence dependence H present, $t = 0$; without H , $t = c$ is delivered independently while a is unchanged. Selling S selects the Muir or VCG mechanism specified below instead of that buyer economy's fixed contracts. Removing H therefore changes an outside resource, not merely a market label.

For each proposed $q > 0$, put $B = b_z + \mu$, $\ell = q/A$, and $w = \theta A$. Supporting household terms are found from the two cash branches,

$$w_{b=0} = \frac{\eta_h[a - B(c - t)]}{1 - (1 + \eta_h)\ell}, \quad (5)$$

$$w_{b>0} = \frac{\eta_h[a - B(c - t) + b_0]}{(1 + \beta_h) - (1 + \beta_h + \eta_h)\ell}. \quad (6)$$

Define disposable surplus $y = a + w\ell - B(c - t)$ and $b = \max\{0, (\beta_h y - b_0)/(1 + \beta_h)\}$; then $x = c - t + (y - b)/B$. Each candidate must satisfy its cash-branch inequality, the procurement bounds, positive log arguments, $x \geq 0$, the household best response and participation. A zero divisor or a nonunique inverse branch requires separate treatment; none is silently included in the declared solutions.

The gate's attained implementation cost is

$$C_{IH}(q) = \min_{\text{feasible supporting terms}} \{\theta q - \mu x + \mathcal{T}(q)\}. \quad (7)$$

For each institutional configuration it maximizes sales receipts minus this cost, with an exit payoff of zero. At $q = 0$, a feasible merchandising-only allocation can instead yield μx ; that option remains distinct from mining. Household budgets imply the consolidated identity $C_{IH} = b_z x + b - a + \mathcal{T}(q)$. The gate's objective is private profit, not social welfare. Every policy optimum is paired with its own zero-charge optimum at identical other primitives.

A.1 Muir receipts and the fixed-route comparator

The separated buyer groups consist of regulated mass $R = 0.005$ with value $M = 2$ and unregulated mass one. Resale opens with probability $\rho_M = 0.02$. Unregulated inverse demand is

$$P_U(u) = \begin{cases} 1 - 8u/3, & 0 \leq u \leq 1/4, \\ 4(1 - u)/9, & 1/4 < u \leq 1. \end{cases} \quad (8)$$

Define $a_M = (6 + \sqrt{6})/48$, $b_M = (1 + \sqrt{6})/8$, $s_M = 1 - 16a_M/3$ and $h_M = 8a_M^2/3$. The concavified unregulated revenue used by the maintained mechanism is

$$v(u) = \begin{cases} h_M + s_M u, & a_M \leq u \leq b_M, \\ u P_U(u), & \text{otherwise.} \end{cases} \quad (9)$$

For the supported charge range $0 \leq \tau \leq 0.10$, the gate receives

$$W_M(q, \tau) = \begin{cases} (M - \tau)q, & 0 \leq q \leq R, \\ (1 - \rho_M)\{v(q - R) + R(M - \tau)\} + \rho_M q P_U(q - R), & R < q \leq 1 + R. \end{cases} \quad (10)$$

This is the declared separated-group illustration, not a claim to solve every Muir environment. Without S , the same gross buyers face fixed routing: a share $\kappa_M = 0.01$ goes to regulated buyers, without resale or reassignment. Receipts are

$$F_M(q, \tau) = \kappa_M q (M - \tau) + (1 - \kappa_M) q P_U((1 - \kappa_M)q), \quad (11)$$

subject to $\kappa_M q \leq R$ and $(1 - \kappa_M)q \leq 1$. Household production additionally requires $q < A$. The base Muir quantity grid contains 160,001 equally spaced candidates on $[0, A]$; infeasible endpoints are discarded and exit remains available.

A.2 Finite whole-buyer VCG and fixed contracts

VCG uses a different buyer economy: buyers $i = 0, 1, 2$ have gross marginal values $g_i(z) = (1.8 - 2z, 0.9 - z, 0.75 - 0.75z)_i$ for $0 \leq z \leq 0.4$. Physical packages have mass $\Delta = 0.000025$. Package j is valued at its midpoint $z_j = (j - 1/2)\Delta$, so its gross bid is $b_{ij} = \Delta g_i(z_j)$. Let $f_i \in \{0, 1\}$ indicate whole-buyer coverage, with $(f_0, f_1, f_2) = (1, 0, 0)$ in the base case. Net cumulative value for k packages is

$$V_i(k; \tau) = \sum_{j=1}^k (b_{ij} - \Delta \tau f_i). \quad (12)$$

For an offered capacity of n packages, calculate

$$\mathcal{W}(n) = \max_{\substack{k_i \in \mathbb{Z}_{\geq 0} \\ \sum_i k_i \leq n}} \sum_i V_i(k_i; \tau), \quad \mathcal{W}_{-i}(n) = \max_{\substack{k_i=0, k_h \in \mathbb{Z}_{\geq 0} \\ \sum_h k_h \leq n}} \sum_{h \neq i} V_h(k_h; \tau). \quad (13)$$

Each buyer's book limits k_i to $0.4/\Delta$ packages. Values are decreasing; ranking positive net marginal bids respects those prefixes. Fixed ties favor the higher buyer index. If $k_i^*(n)$ is the selected allocation, the whole-buyer externality payment is

$$T_i(n) = \mathcal{W}_{-i}(n) - [\mathcal{W}(n) - V_i(k_i^*(n); \tau)]. \quad (14)$$

For produced stock $q = N\Delta$, the gate can offer any integer $0 \leq n \leq N$. Unsold stock has zero retention value. It chooses the highest available $\sum_i T_i(n)$ and then chooses N against the attained household cost. Thus offered capacity and extracted stock are separately optimized; the policy charge shifts covered buyers' bids rather than being deducted a second time from the gate's receipts.

Without S , predetermined quotas have base shares $(0.25, 0.50, 0.25)$. Integer packages are assigned by largest remainders, with buyer index breaking equal remainders. Buyer i pays its lowest allocated gross marginal value per mineral unit, less τf_i , on its entire allotment. A negative price makes that candidate infeasible. A change in the fixed share of buyer 0 redistributes the remainder two-thirds to buyer 1 and one-third to buyer 2; it does not change coverage flags. A coverage stress instead changes the whole covered buyer.

A.3 Declared sensitivity and finite scope

The five added joint bundles below each use charges 0.05 and 0.10, paired with zero. Unlisted primitives retain their base values. The notation in the last column is (Muir fixed regulated share, VCG fixed buyer-0 share, covered VCG buyer).

The reproduced source set additionally changes one primitive at a time: halves the charge, sets wealth to 0.01 or 0.03, halves the margin cap, doubles the Muir fixed share and the VCG buyer-0 fixed share, or covers buyer 1 in VCG. Its refinement doubles Muir quantity resolution and goods-margin resolution and halves VCG package mass. Every optimum is certified only over its stated finite quantity, margin and capacity alternatives, with the exact household branches

Table 8: Added joint primitive bundles

Bundle	$(s_o, a, c, \bar{\mu})$	$(\kappa_M, s_0, i_{\text{covered}})$
Low outside option, high wealth	(.004, .030, .040, .25)	(.01, .25, 0)
High outside option, low wealth	(.016, .010, .060, .50)	(.01, .25, 0)
No margin, high fixed exposure	(.008, .020, .050, 0)	(.02, .50, 0)
Tighter participation, losing buyer covered	(.012, .020, .060, .50)	(.01, .25, 1)
High outside option, exit stress	(.032, .010, .060, .25)	(.02, .50, 1)

above. There is no empirical weighting of these bundles and no claim that finite-grid endpoints bound all intermediate primitives.

B Incidence, institutional comparisons and aggregation

Let $Q = O + E$ with candidate extraction Q , outside sales O and candidate certificate entries E . Maintain $pO = D > 0$ and $Q = Bp^\eta$, with $B, \eta > 0$, between compared screening states. For $\delta = 1 - Q_1/Q_0 \in (0, 1)$,

$$\frac{p_1}{p_0} = (1 - \delta)^{1/\eta}, \quad \frac{O_1}{O_0} = (1 - \delta)^{-1/\eta}, \quad (15)$$

$$E_0 - E_1 = (Q_0 - Q_1) + (O_1 - O_0) \quad (16)$$

$$= Q_0\delta + O_0[(1 - \delta)^{-1/\eta} - 1]. \quad (17)$$

All terms in the last line are nonnegative and $Q_0\delta > 0$, so the denominator is strictly positive. Thus

$$T = \frac{Q_0 - Q_1}{E_0 - E_1} = \frac{\delta}{\delta + (O_0/Q_0)[(1 - \delta)^{-1/\eta} - 1]}. \quad (18)$$

Since the derivative of $(1 - \delta)^{-1/\eta}$ at zero is $1/\eta$, the local limit is $T_{\text{loc}} = \eta/(\eta + O_0/Q_0)$. The accounting identity $\delta = T(E_0 - E_1)/Q_0$ explains why high incidence of a small displaced flow can coexist with a small aggregate contraction.

The procurement identity $\zeta = \eta/(1 + \eta) + m$ gives $\eta = (\zeta - m)/(1 - \zeta + m)$ on $0 < \zeta - m < 1$. For fixed $o = O_0/Q_0 > 0$, $d[\eta/(\eta + o)]/d\eta = o/(\eta + o)^2 > 0$. Under $0 \leq m < \zeta < 1$, the local incidence has limiting lower bound zero and upper value at $m = 0$. If negative m is admitted, η need not have a finite upper bound; incidence still lies below one and approaches one as $\eta \rightarrow \infty$. These are conditional local restrictions, not empirical uncertainty intervals.

For channels $N = \{S, I, H\}$ and configuration $J \subseteq N$, define $R(J) = Q(0, J) - Q(\tau, J)$. Every configuration uses its own no-charge allocation. Absorbed quantity is $A = R(\emptyset) - R(N)$; its fraction is $A/R(\emptyset)$ only when that denominator is meaningful. The Shapley contribution is

$$\varphi_i = \sum_{J \subseteq N \setminus \{i\}} \frac{|J|!(2 - |J|)!}{3!} [R(J) - R(J \cup \{i\})]. \quad (19)$$

For each ordering of channels, incremental changes telescope from $R(\emptyset)$ to $R(N)$. Averaging over the six orders gives $\sum_i \varphi_i = A$. A negative component is permitted and must not be discarded.

The household budget is $(b_z + \mu)x + b = a + \theta q$. The gate's private implementation account therefore satisfies

$$\theta q - \mu x + \mathcal{T}(q) = b_z x + b - a + \mathcal{T}(q). \quad (20)$$

Here x is the household basket, b retained cash, a outside wealth, b_z the basket resource price, μ the goods margin and θ procurement pay per unit. This x is unrelated to the screening model's price ratio. The identity reconciles recovery with resources; it is not a welfare aggregator.

On the supported Muir branch $q > R$, receipts have the form

$$W(q, \tau) = (1 - \rho)[v(q - R) + R(M - \tau)] + \rho q P_U(q - R). \quad (21)$$

Holding its declared primitives fixed, $W_\tau = -(1 - \rho)R$ and $W_{q\tau} = 0$. By contrast fixed-route sales have $W_\tau = -\kappa q$ and $W_{q\tau} = -\kappa$. The first invariance holds only while the supported branch and operation remain optimal. Finite VCG payments are reconstructed from each buyer's externality, including recalculation with that buyer removed; no first-order condition substitutes for the finite allocation comparison.

For recorded destination shares s_{it} and positive origin weights w_{it} summing to one in the same comparable cohort,

$$\Delta \sum_i w_i s_i = \sum_i \frac{w_{i1} + w_{i0}}{2} \Delta s_i + \sum_i \frac{s_{i1} + s_{i0}}{2} \Delta w_i. \quad (22)$$

Expanding each product cancels the cross terms exactly. A separate full-population/cohort bridge is necessary when available origins differ.

C Independent checks and the evidence gate

The verifier reconstructs all 560 allocations without importing the generating optimizer, enumerates the declared finite quantity and goods-margin choices, and recomputes household conditions, Muir receipts, VCG allocations/payments and all 35 Shapley decompositions. The largest discrepancy in finite-grid best profit is below 2.8×10^{-17} ; the household first-order residual is below 2.7×10^{-15} . Deliberately corrupting implementation cost and a Shapley component produces the expected failures. These results establish finite-grid optimality and identities, not a continuous global optimum or empirical identification.

The archive contains 2,426 underlying trade rows for the verified transformations and 15 empty empirical templates. The rows do not supply policy assignment, and the empty templates are not a sample. Historical forecast windows are already-used diagnostics. Consequently the new model effects have no licensed statistical confidence interval or p-value. An executable acquisition validator and the prospective design accompany the package; estimation remains gated on actual observations and a defensible design.

D Claim scope and replication record

The accompanying certificate distinguishes five questions. Numerical consistency concerns the executed equilibria and accounts. Measurement support concerns the source's actual definitions and coverage. Identification concerns which model quantities the observations determine. External validity concerns the population and institutions to which the result can travel. Policy appraisal additionally requires implementation costs, budgets and outcome evidence. Passing the first question does not imply the others.

The result is certified conditional on the specified primitives, classifications, equilibrium branches and finite experiments. The stronger proposition that these numbers identify the effect

of a named implemented policy requires modification. Statistical policy-effect inference is NOT ESTIMABLE FROM CURRENT EVIDENCE. All ranges preserve this status; missing confidence levels and p-values remain null rather than zero.

The replication package contains the frozen specification, source-admission and identification ledgers, scenario-level effects, primary derivations, independent first-stage reconstructions, adversarial findings and a clean-run entry point. Source and result hashes preserve provenance. Original manuscripts and baseline outputs are not overwritten. The next appendix records the complete common execution prompt so that the same evidentiary constraints govern all three contributions.

E Country scales, assumptions and missing calibration

The following registry is shared by the structural exercises. “Range” means the assumed candidate fractions 0.1, 0.5 and 0.9, not measured illegal or artisanal shares. Existing gold and DRC cobalt fractions retain their original transport qualifications. ROW is the source residual with a fixed 0.5 candidate split, no vulnerable certificate origin and no government. Uganda tantalum and DRC/Burundi/Uganda tungsten remain descriptive BGS cases without a named structural scale. Zero production and missing calibration are distinct. The fuller BGS producer tabulation is supplied as a separate CSV source scenario; its country coverage should not be confused with the narrower gold calibration.

Table 9: Admitted structural production registry

Mineral	Country or residual	Tonnes	Candidate share
Gold	Peru	108.15	0.1873 (transported)
Gold	Colombia	56.70	0.5386 (transported)
Gold	Rest of world and source rounding balance	3,496.35	0.5 (assumed residual)
Cobalt	United States	200.00	Range (assumed)
Cobalt	Australia	4,780.00	Range (assumed)
Cobalt	Canada	3,350.00	Range (assumed)
Cobalt	China	2,000.00	Range (assumed)
Cobalt	Democratic Republic of the Congo	226,000.00	0.1000 (transported)
Cobalt	Cuba	3,450.00	Range (assumed)
Cobalt	Indonesia	35,000.00	Range (assumed)
Cobalt	Madagascar	3,100.00	Range (assumed)
Cobalt	Papua New Guinea	2,630.00	Range (assumed)
Cobalt	Philippines	3,100.00	Range (assumed)
Cobalt	Russia	8,000.00	Range (assumed)
Cobalt	Turkey	2,200.00	Range (assumed)
Cobalt	Rest of world and source rounding balance	8,190.00	0.5 (assumed residual)
Tantalum	United States	0.00	Zero scale; no experiment
Tantalum	Australia	52.00	Range (assumed)
Tantalum	Bolivia	2.00	Range (assumed)
Tantalum	Brazil	210.00	Range (assumed)
Tantalum	Burundi	2.00	Range (assumed)
Tantalum	China	76.00	Range (assumed)
Tantalum	Democratic Republic of the Congo	1,270.00	Range (assumed)
Tantalum	Ethiopia	40.00	Range (assumed)
Tantalum	Mozambique	55.00	Range (assumed)
Tantalum	Nigeria	390.00	Range (assumed)
Tantalum	Russia	29.00	Range (assumed)
Tantalum	Rwanda	374.00	Range (assumed)

Mineral	Country or residual	Tonnes	Candidate share
Tin	United States	0.00	Zero scale; no experiment
Tin	Australia	11,300.00	Range (assumed)
Tin	Bolivia	21,200.00	Range (assumed)
Tin	Brazil	27,600.00	Range (assumed)
Tin	Myanmar	20,000.00	Range (assumed)
Tin	China	71,000.00	Range (assumed)
Tin	Democratic Republic of the Congo	26,000.00	Range (assumed)
Tin	Indonesia	55,000.00	Range (assumed)
Tin	Laos	1,860.00	Range (assumed)
Tin	Malaysia	5,460.00	Range (assumed)
Tin	Nigeria	3,100.00	Range (assumed)
Tin	Peru	32,300.00	Range (assumed)
Tin	Russia	3,260.00	Range (assumed)
Tin	Rwanda	4,100.00	Range (assumed)
Tin	Vietnam	11,000.00	Range (assumed)
Tin	Rest of world and source rounding balance	820.00	0.5 (assumed residual)
Tungsten	United States	0.00	Zero scale; no experiment
Tungsten	Australia	920.00	Range (assumed)
Tungsten	Austria	840.00	Range (assumed)
Tungsten	Bolivia	1,700.00	Range (assumed)
Tungsten	China	67,000.00	Range (assumed)
Tungsten	Kazakhstan	0.00	Zero scale; no experiment
Tungsten	North Korea	1,900.00	Range (assumed)
Tungsten	Portugal	650.00	Range (assumed)
Tungsten	Russia	1,500.00	Range (assumed)
Tungsten	Rwanda	1,300.00	Range (assumed)
Tungsten	Spain	700.00	Range (assumed)
Tungsten	Vietnam	3,400.00	Range (assumed)
Tungsten	Rest of world and source rounding balance	2,090.00	0.5 (assumed residual)

F Measurement and procurement admission

The oxide conversion for tantalum uses the declared Ta fraction 0.81897 in Ta_2O_5 . A quotation of USD 167 per kg oxide-equivalent therefore gives $167000/0.81897$ USD per tonne of contained Ta. Tungsten uses the declared 0.793 W fraction in WO_3 : one dry metric tonne unit is 10 kg WO_3 , or 7.93 kg W, giving $252000/7.93$ USD/t W for a USD 252/dmtu reference. Cobalt and tin pound quotations use 0.45359237 kg/lb. All these conversions are checked separately from their suitability as procurement prices.

The Rwanda 3T study provides evidence of assays, payment stages and food deductions, but its tantalum table is not admitted as an elasticity moment. The COMIKAGI table reports 1,540 kg concentrate per month at 30% Ta_2O_5 , a USD 130/kg quotation of unclear basis, and USD 37.7/kg concentrate locally. The displayed international total USD 200,200 treats the quotation as per kg concentrate; an oxide-content application instead gives USD 60,060. The table’s combined total USD 514,578 also differs from USD 45,619 plus USD 200,200. The local total USD 58,058 is internally consistent. These alternatives imply radically different ratios, so no tantalum procurement elasticity is inferred. ([Rwanda study, 2018](#))

The cassiterite example is arithmetically coherent: 3,960 kg concentrate, 60% tin and USD 19.20/kg metal imply approximately USD 45,619 internationally; USD 11/kg concentrate implies USD 43,560 locally. The ratio is approximately 0.955. Gifurwe’s tungsten comparison gives approximately 0.686 from USD 8.4/kg concentrate and USD 188.5/dmtu at 65% WO_3 . Both are transaction-stage ratios, not supply elasticities. The independently inspected pages and all arithmetic are retained in the Rwanda procurement audit.

Observed scales, conditional fits and assumptions have distinct records. No new mineral inherits the Colombia or Peru procurement elasticity merely because it is artisanal or traded through the same region. Identifying that elasticity requires matched transaction prices, assay and deduction data, production costs, net margins and a defensible population bridge.

G Reproduction and limits of verification

The original thirteen reproduction stages pass before extension. The archive preserves original source hashes and a runnable benchmark snapshot. Resource configurations validate names, units and sources; an unknown mineral raises an error. Every saved effect joins a model, mechanism, mineral, country, source scenario, policy, comparator, horizon and outcome. An unavailable pair is not assigned a zero effect.

The numerical admission targets are 10^{-8} for static and research accounts and 10^{-6} for government or Bellman conditions, with separate finite-grid and horizon comparisons. An economic-domain rejection and numerical nonconvergence are different outcomes. Some expanded-country government iterations cycle near indifference. Such a case may be admitted only when a separately reconstructed full-country equilibrium satisfies the original economic-equation thresholds; its fixed-point iteration flag remains false and its distinct status is retained. This does not establish exact-grid uniqueness or stable intervention timing.

All reported ranges are finite conditional exercises. Related model variants share theoretical structure and evidence; their agreement is not independent empirical confirmation. Policy-effect confidence intervals and p-values remain unavailable without an identifying dataset and design. Government financing needs remain country-specific; no international transfer arrangement is inferred from aggregation. A reference-value index never adds tonnes of different metals.

References

United Nations Statistics Division. UN Comtrade Database. Archived 2026 extraction of partner-import records; exact product, query, reporter coverage and response hashes in the accompanying replication archive. [Primary source](#).

UNODC and Ministerio de Minas y Energía. 2025. Colombia: Explotación de oro de aluvión. Evidencias a partir de percepción remota 2023. Procurement passage p. 50; production categories Table 5, p. 56. [Primary source](#).

- Banco Central de Reserva del Perú. 2026. Memoria 2025, commercialization passage, indexed chapter. Provisional: archived nominal PDF is HTML; underlying USAID observation date and sample unresolved. [Primary source](#).
- PRG, IPIS, SFR, and Ulula. 2020. Evaluating Due Diligence Programs for Conflict Minerals: A Matched Analysis of 3T Mines in Eastern DRC. November, Table 12, p. 44. [Primary source](#).
- U.S. Geological Survey (2026). *Mineral Commodity Summaries 2026*: cobalt, tantalum, tin and tungsten country tables and price notes. <https://www.usgs.gov/centers/national-minerals-information-center/commodity-statistics-and-information>. Archived source vintage and pages in the evidence ledger.
- British Geological Survey (2026). *World Mineral Statistics*, country production series and source notes, 2005–2024 extraction. <https://www.bgs.ac.uk/mineralsuk/statistics/world-mineral-statistics/>.
- United Nations (2026). *UN Comtrade*, archived annual origin and destination queries and selected bilateral 3T routes. <https://comtradeplus.un.org/>.
- United Nations Statistics Division (2017). Harmonized System code 261590: niobium, tantalum and vanadium ores and concentrates. <https://unstats.un.org/unsd/classifications/Econ/Detail/EN/32/261590>.
- International Peace Information Service. DRC and CAR mine-visit datasets, dated observations, public ODC-BY data; archive retrieved 2026. <https://ipisresearch.be/home/maps-data/maps-of-drc/interactive-map-of-artisanal-and-small-scale-mining-exploitation-in-eastern-drc/>.
- Extractive Industries Transparency Initiative. Country summary workbooks and reports, archived exposed vintages; country-year records documented individually. <https://eiti.org/>.
- Kimberley Process. Rough diamond production, trade, value and certificate statistics; dated country panel. https://kimberleyprocessstatistics.org/public_statistics.
- Economic Contributions of Artisanal and Small-Scale Mining in Rwanda: Tin, Tantalum, and Tungsten*. COMIKAGI and Gifurwe case-study tables, printed pp. 60–61 and 75–76. https://assets.publishing.service.gov.uk/media/5a392cf4e5274a79051c9d82/Rwanda_case_study.pdf. Procurement arithmetic audited separately.