

Certification, procurement and production

What a cleaner register can establish

Elies Lucio

Calibration 1 · 4 October 2026

Abstract

This global edition extends the preserved gold and cobalt calibration to tantalum, tin and tungsten in separate world markets. The original thirteen reproduction stages pass. A 3,780-case static design admits 3,332 cases, with economic-domain exclusions retained. Under common central assumptions, stronger discovery changes live certificate composition far more than extraction; outside spending remains a separate determinant of physical production. Supply exponent one supplies the curvature bridge to alternative nonlinear VCG and Muir economies. A shared production, trade and institutional archive distinguishes observed scales, transported moments and assumed strata. An audit of the Rwanda tantalum table withholds its ambiguous procurement ratio. Detailed country scales and conditional effects accompany the global comparison. These are model sensitivity results; available evidence does not identify causal policy effects, confidence intervals or p-values.

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 The institutional problem

A buyer who wishes to avoid disallowed origin confronts two decisions. The first concerns which claims deserve recognition; the second concerns how much to pay for material that qualifies. These decisions affect the seller’s return to obtaining certification. They need not remove the seller’s ability to supply another buyer. A more reliable register can therefore improve the composition of certified trade even when the outside market continues to sustain extraction.

That possibility is economically important because certification success is often described using a fraction: the share of live certificates attached to candidate material. The numerator is a quantity of material, while the denominator includes permitted material entering the register. The fraction falls when fewer candidate claims survive discovery. It also falls when permitted

entries expand. An evaluation of extraction must follow the physical numerator outside the register as well as within it.

Here candidate means the stratum to which the model applies its provenance concern. It does not mean that every observed small producer, artisanal miner or barequero is illegal, controlled by an armed actor or environmentally harmful. The available administrative categories establish useful scales and exposure scenarios; they do not identify those substantive conditions. This distinction matters for distributional interpretation. A loss of receipts in a candidate stratum is a modeled incidence on that stratum, not an estimate of a reduction in criminal rents or a measured household welfare cost.

Institutions shape the returns to production and the distribution of those returns ([Acemoglu, Johnson, and Robinson, 2005](#)). The experiment therefore asks a limited but informative question: under stated supply and expenditure assumptions, how much of improved certification reaches candidate production? It separately measures the consequences of changing the outside market's spending. Peru and Colombia share one gold market, so a policy affecting either origin can alter the price faced by both. Cobalt clears in its own market. The original CES exercise is retained as a comparison with a different demand and certification mechanism, rather than treated as another estimate of the same intervention.

2 Incentives, transactions and measurement

The contractual mechanism distinguishes genuine permitted production, candidate production, entries in the register and material sold outside it. A vulnerable origin can support an inflated claim by diverting some candidate material into a permitted account. Fraud and diligence respond to the contract. Discovery removes some diverted material from live certificates, and the remainder stays in the live register. The origin must possess enough candidate material to back the diversion. A numerical solution that fails that stock requirement is economically inadmissible even if aggregate spending equations clear.

The reference contract uses late discovery of 0.25, a bond parameter of 0.20, forgery cost of 2, diligence curvature of 0.5 and a fee of 2 percent. These are assumptions. The new contract grid raises discovery to 0.35 or 0.50 and either retains diligence curvature at 0.5 or halves it to 0.25. Discovery is a model probability; none of the admitted administrative or trade records estimates it. The reference contract and policy are evaluated under the same supply population and spending whenever their difference is called the contract effect.

Physical scale is attached to 2024 observations. Peru contributes 108.148345 tonnes of recorded fine gold; the selected small and artisanal categories sum to 20.256299 tonnes. Colombia contributes 56.7 tonnes in the frozen royalty-associated vintage. Applying the separate 2023 barequero share to that figure produces a 30.539416-tonne candidate scenario. The share is calculated with the numerator and denominator from the same source table; its transport to 2024 is an assumption. DRC cobalt uses the revised 2024 vector of 226,000 tonnes nationally and 302,000 tonnes worldwide. The 10 percent candidate share is a historical ASM scale transported to 2024, yielding 22,600 scenario tonnes. It is not a measured contemporary illegal share. ([MINEM, 2025](#); [ANM–UPME, 2025](#); [UNODC and MME, 2025](#); [USGS, 2026](#); [Gulley, 2023](#))

The world gold mining scale is 3,661.2 tonnes. The rest-of-world split and the share of permitted output vulnerable to false claims are chosen. Consequently, world candidate output has a large assumed component. Country and world effects are reported separately. Monetary calculations use nominal reference-price units: the gold price is a fine-metal benchmark, and the

cobalt price is a refined-metal benchmark. Candidate gross receipts are price times candidate output at the model’s outside quote. They are neither realized minegate cash nor disposable household income. ([World Gold Council, 2025](#))

Procurement evidence is informative about a transaction stage, but does not directly measure supply response. Colombia’s reported purchase-to-LBMA ratio of 0.70 concerns selected field transactions. Exact dates, sampling uncertainty, assay adjustments and deductions are unavailable in the cited passage. Peru’s reported 0.85–0.90 range concerns regional amalgamated-gold procurement and remains provisional: the accessible indexed passage and the archived HTML file do not constitute a successfully retrieved PDF or a transaction dataset. Applying either proxy to the entire national candidate stratum requires a separate transport assumption. ([UNODC and MME, 2025](#); [BCRP, 2026](#))

Under a power-cost implementation, the cost share at the optimal quantity is $\eta/(1+\eta)$. If the procurement ratio equals that cost share plus a net margin m , the bridge is $\eta = (\zeta - m)/(1 - \zeta + m)$. Setting the margin to zero produces Colombia’s familiar 7/3 and Peru’s 17/3 to 9. Those numbers are consequences of a zero-margin restriction. In the chosen margin grid of 0, 0.10 and 0.20, Colombia’s conditional elasticity ranges from 1 to 7/3; Peru’s from approximately 1.857 to 9. Without a restriction on the margin and a matched price-stage bridge, the observed ratio alone does not identify η .

3 A paired experiment with explicit supply alternatives

The common-curvature branch considers supply curvature η equal to 0.25, 0.5 or 1. The heterogeneous gold branch applies the procurement bridge only to the selected Colombian and Peruvian candidate supply. Other origins and permitted supply retain the chosen common curvature. This prevents a Colombian field ratio from silently determining the technology of Peru, legal industry and the rest of the world. It still assumes that the regional procurement relation can represent the selected national stratum, and it deliberately exposes that assumption through the grid.

Each supply branch crosses the contract grid with initial outside-to-certified quotes of 0.70, 0.80 and 0.85 and outside-spending multipliers of 0.80, 1 and 1.20. The physical baseline is refitted at every initial quote; observed country quantities are not allowed to drift merely because a normalization changes. The spending stress is selected, not estimated. Every valid case yields three paired differences: contract versus reference at the same spending, spending alone versus the anchored baseline, and their combined change. These contrasts add exactly in levels. Their percentage changes have their own comparator denominators and should not be added.

Gold is solved both with mining alone and with 1,370 tonnes of fixed recycled supply. The alternative treats the recycled flow as fully permitted and adds its demand backing at the baseline. It assigns recycling no price elasticity. The resulting equilibrium is cleared again, rather than multiplying a mining result by a recycling share. Above-ground stocks, price-responsive scrap collection and inventory demand are outside this minimal closure. ([World Gold Council, 2025](#))

Cobalt is solved as a single metal and under fixed-mix joint revenue. The latter uses the industrial comparator weight $8.743/(41.857 + 8.743)$, or 0.172787, for cobalt’s share of copper-cobalt revenue. Supply responds to the index $1 - w + w(p/p_0)$, raised to the chosen joint-revenue curvature. Other-metal prices and physical mix are fixed. This is an explicit transport of an industrial accounting ratio into a hypothetical supply closure, not a finding about artisanal

revenue composition. Local cobalt-price elasticity at the baseline is η times w , but both the joint technology and its application to the modeled population are assumed. (CMOC, 2025)

The scheduled grid contains 2,376 cases: 2,160 gold cases and 216 cobalt cases. Of these, 2,074 clear within the maintained domain and 302 exceed a necessary upper bound on outside spending in the strict positive-fraud branch. None remains numerically unsolved. The rejected cases are retained in the denominator, with the reason recorded. A corner with zero fraud would require a separately specified equilibrium; it is not silently substituted. The numerical ranges below concern only the valid cases in this chosen grid and are explored extrema, not proved bounds over continuous parameter regions.

4 Results: large register effects, small direct extraction effects

At the reference physical calibration, raising late discovery from 0.25 to 0.50 reduces world candidate extraction by 0.005913 percent for gold and 0.286023 percent for cobalt. Surviving candidate mass in live certificates falls by 57.75 and 55.51 percent, respectively. The gap is a feature of the mechanism: changing the return to misrepresenting origin substantially changes the composition of certified entries, while outside spending continues to support almost the same physical supply.

The expanded grid preserves that distinction. At fixed outside spending, common-curvature gold contraction ranges from 0.00120 to 0.01513 percent. In the heterogeneous mining-only gold branch, the Peruvian candidate stratum contracts by 0.00561 to 0.21164 percent, or 0.00129 to 0.06492 tonnes in the matched comparisons. Colombia contracts by 0.00301 to 0.05491 percent, or 0.00101 to 0.01252 tonnes. Percentage and absolute extrema can arise in different cases because spending changes the comparator’s scale.

The corresponding DRC candidate contraction is 0.05728 to 0.74256 percent in the single-metal branch, or 13.37 to 151.75 tonnes. Joint-revenue support reduces it to 0.01290–0.19101 percent, or 2.94–42.01 tonnes. The fixed recycling alternative changes the gold results only slightly in this design. These comparisons do not establish that gold recycling is generally unimportant. They establish that a fixed permitted flow, with its baseline spending backing, has little influence on the particular contract transmission considered here.

Table 1: Extraction responses under matched and changing outside spending

Candidate stratum / closure	Contract effect (%)	Combined effect (%)
Peru, regional curvature	-0.21164 to -0.00561	-77.82 to 205.59
Colombia, regional curvature	-0.05491 to -0.00301	-32.56 to 36.05
DRC, single metal	-0.74256 to -0.05728	-10.25 to 8.68
DRC, joint revenue	-0.19101 to -0.01290	-2.87 to 2.49

Notes: Policy discovery is 0.35 or 0.50 versus 0.25, crossed with the declared diligence grid. Contract effects hold outside spending fixed within each pair. Combined effects include a spending multiplier of 0.80, 1 or 1.20. Entries are explored extrema among admissible cases; no statistical inference is available.

Live candidate certificate mass responds much more. Across valid gold procurement cases, it falls by roughly 22.25–98.99 percent in mining-only markets; the fixed recycling alternative reaches approximately 99.14 percent at an explored endpoint. Cobalt’s live candidate mass falls by 21.68–84.12 percent as a single metal and 21.57–74.84 percent under joint revenue. A register objective and an extraction objective can therefore rank a policy’s apparent success quite differently. Neither measure alone describes the displacement of sellers or environmental harm.

Outside spending dominates the sign of the combined extraction contrast. In heterogeneous gold mining, the Peruvian range is -77.82 to +205.59 percent and the Colombian range is -32.56 to +36.05 percent. The large Peruvian expansion is possible because high conditional curvature is applied to a relatively small baseline stratum in a shared world market. It is an exposed consequence of the procurement and margin restrictions, not evidence for a threefold real-world response. The DRC combined range is -10.25 to +8.68 percent as a single metal and -2.87 to +2.49 percent under joint revenue.




Figure 1: Outside spending changes the scale and sign of the comparison

Notes: Separate horizontal scales. Horizontal segments connect explored extrema, not confidence limits. Gold uses heterogeneous procurement-curvature assumptions; cobalt compares single-metal and industrially transported joint-revenue closures. Negative changes denote contraction.

An illustrative target of 5 percent candidate contraction is reached by no contract-only case. With the selected spending stresses, some combined gold and single-metal cobalt cases pass that target, while none of the joint-revenue cobalt cases does. No official target or welfare threshold has been supplied, so five percent is a transparent illustration. The useful general object is the spending frontier: the increase in outside spending at which candidate production returns to its baseline under the stronger contract. Across the chosen cases, it is only 0.00690–0.02348 percent for gold and 0.361–1.286 percent for cobalt. These small thresholds identify the demand-transmission question that an empirical evaluation must resolve.

Receipts need separate interpretation. In the heterogeneous mining-only branch, contract-only modeled candidate gross receipts fall by 0.00863–0.23513 percent in Peru and 0.00602–0.07844 percent in Colombia. The single-metal DRC range is 0.279–1.480 percent; under joint revenue, cobalt receipts fall by 0.276–1.462 percent even though cobalt output is more stable. Copper revenue is excluded from that cobalt-receipt measure. Fee collections are recorded, but incremental discovery costs are not observed. A fee ledger therefore cannot establish the affordability of stronger enforcement.

5 Innovation, denominators and financial backing

The original physical model’s innovation and collateral calculations were freshly reproduced for periods 0–100, with its spending-based research calibration retained as a conditional assumption. At period 60, the gold leakage share among live certificates declines from 0.034695 to 0.013598

percent. Yet candidate extraction is 0.000285 percent above its initial level and candidate mass in live claims is 1.671 percent higher. For cobalt, the leakage share declines from 0.475217 to 0.243862 percent, while extraction rises 0.187028 percent and live candidate mass rises 20.256 percent. (OECD, 2025)

These are changes along an imposed innovation and expenditure trajectory relative to the initial state. They are not separately identified effects of a real intervention. Their value is diagnostic: expanding permitted capacity and certified expenditure can dilute the leakage ratio while leaving physical candidate activity intact. The baseline research spending ratio does not identify the assumed innovation technology, and a model period has no calibrated calendar duration.

Collateral likewise requires more than a positive fee balance. Refundable principal belongs to the trust ledger; it is not available service revenue. Matured deposits split into refunds and forfeitures. Service cash, research grants, fee receipts and the opportunity cost of immobilized principal have different timing and incidence. At period 60, the reproduced central service cash gap is zero in both markets. Including the assumed 5 percent per-model-period capital opportunity charge produces economic funding gaps of approximately USD 6.402 billion for gold and USD 271.305 million for cobalt in nominal benchmark-scaled units. Outstanding principal is approximately USD 219.573 billion and USD 9.107 billion. These magnitudes describe the assumed world-scale contract, not audited agency budgets or required appropriations in a particular country.

An independent reconstruction follows each deposit vintage, maturation, forfeiture and refund across all 202 market-dates. It also checks the timing of grant commitments and payments and recomputes the two funding-gap definitions. Together with origin backing and market clearing, this establishes accounting consistency for the executed path. It does not establish that participants would provide the collateral, that administrative capacity exists, or that the assumed capital charge reflects an observed social cost.

6 What the evidence can support

The original CES comparison illustrates the dependence on institutional closure. Its freshly reproduced reference candidate contractions are 4.2595 percent for gold and 3.9597 percent for cobalt, with wide sign-changing extrema on its own assumption grid. Those figures refer to a demand segmentation experiment and retain the original cobalt vintage. They are not pooled with the contracting results, which use the revised physical vector and a distinct intervention. (USGS, 2025)

For every effect reported here, statistical significance is **NOT ESTIMABLE FROM CURRENT EVIDENCE**. Country administrative totals, reported procurement ratios and corporate accounts provide scales or conditional comparisons. They do not supply independent treatment units, a sampling frame for the relevant transactions, valid assignment variation, or an identified response equation. The number of grid cases is not an empirical sample size. Calling the results statistically insignificant would be equally misleading: no valid zero-effect test has been performed.

The missing evidence has a concrete acquisition design. Repeated lot-level observations should connect mine or producing organization, fine-metal assay, transaction date, gross quote, transport and finance deductions, and alternative buyers. The same lots should link entry, audit, discovery, revocation and final sale. A phased or otherwise credibly assigned enforcement change would

need known assignment units, dates and spillover exposure. Household outcomes require their own linked panel of earnings, employment and consumption; national receipts cannot substitute for it. Copper credits, cobalt grades and joint production costs must be observed before the industrial revenue bridge can become an artisanal supply estimate. Before estimation, the design must freeze outcomes, assignment and observational units, clustering, missingness, treatment reach and the confirmatory hypothesis family.

The contribution is therefore a conditional statement about incentives and measurement. Stronger certification can sharply reduce the candidate material represented by surviving claims. The extent to which it also reduces extraction is governed by outside purchasing power, production response and the institutional ability to keep displaced material from finding an equally remunerative outlet. Those links require evidence distinct from a cleaner register.

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 2: 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 extension: certification across five mineral markets

The extension retains the contract and supply equations above and replaces commodity-specific branches with a validated resource registry. It executes 3,780 scheduled scenarios: 3,332 satisfy the maintained economic domain, 448 are explicitly rejected and none remains numerically unsolved. The full grid preserves exponents 0.25, 0.5 and 1; initial outside-to-certified quotes 0.70, 0.80 and 0.85; outside spending 0.80, 1 and 1.20; and the original discovery and diligence alternatives. The procurement-derived gold, fixed-recycling gold and fixed-mix cobalt alternatives remain separately named closures.

The common five-mineral comparison fixes the initial quote ratio at 0.8 and raises discovery from 0.25 to 0.50 at unchanged outside spending. Unmeasured country candidate fractions are 0.5 in this displayed scenario and range from 0.1 to 0.9 elsewhere. The registry fixes all observed physical scales while changing those fractions. The enlarged cobalt country composition therefore differs from the reproduced DRC-plus-residual benchmark. That compositional change explains why its new central result need not equal the older cobalt figure in the benchmark discussion.

Table 3: Matched global screening comparisons

Mineral	Extraction, $\eta = .5$	Extraction, $\eta = 1$	Live claims, $\eta = .5$
Gold	-0.00591	-0.00888	-57.75
Cobalt	-0.32895	-0.50559	-55.18
Tantalum	-0.08937	-0.13541	-56.72
Tin	-0.08913	-0.13504	-56.72
Tungsten	-0.08719	-0.13207	-56.75

Notes: Percentage changes in candidate quantities. Each policy uses its own matched mineral baseline. These are world totals under assumed strata, not country estimates of illegal mining.

For the three new metals, the shared normalized assumptions produce similarly small direct extraction responses and large register-composition changes. This similarity is a consequence of the maintained technology, outside demand and source composition. It is not three new empirical estimates. The greater value of the extension is to expose which country scales and assumptions would have to change for a different production response.




Figure 2: Certification composition and extraction respond at very different scales

Notes: Independent horizontal scales preserve the difference in magnitudes. The plotted central scenario is defined in the accompanying table.

Spending-only, screening-at-common-spending and combined effects are reconstructed separately. Their changes add in levels, not by adding percentages with different denominators. The separate frontier exercise contains 2,520 scheduled zero- or 5%-contraction targets, of which 2,484 have a licensed solution in the stated branch; the maximum target residual is 2.86×10^{-14} . A required reduction in buyer spending is not an estimate of the public cost of achieving it.





Figure 3: Production effects depend on the outside-spending assumption

9 Comparison with nonlinear VCG and Muir economies

The alternative papers in `consolidated calibration` use the same observed reference bundle and quadratic extraction costs, corresponding to $\eta = 1$. The second extraction column is therefore the curvature bridge. Each nonlinear mechanism has its own calibrated baseline, whole-buyer or continuum allocation, retention choice and final physical output. Comparing either to the old $\eta = .5$ run without this bridge would confound changed technology and institutions.

The shared conclusion is conditional: a certificate measure need not track physical extraction, and outside spending remains a separate transmission margin. Mechanism-specific payment, resale and retention accounts determine the size of that divergence. The models cannot be ranked as empirically superior with these aggregate observations alone. The paired synthesis records numerical differences, common assumptions and source-dependent alternatives explicitly.

Worldwide BGS production records, 2024 | Part 1




Export proxies (hatched): Gold: Togo, Rwanda, Madagascar, Ethiopia. Tin: Mongolia, Rwanda.

BGS is a separate source scenario, not the USGS calibration. Missing is not zero.
Ta uses contained-metal rows; diamonds remain carats. Country notes retain all flags.

Sources: BGS API (4 October 2026); Natural Earth. See the accompanying country table.

Figure 4: Global production in the separate BGS descriptive source

Notes: Source definitions and export-proxy flags remain explicit. Gray is unavailable in the plotted population, not necessarily zero. Each mineral has its own source denominator.



BGS is a separate source scenario, not the USGS calibration. Missing is not zero.
 Ta uses contained-metal rows; diamonds remain carats. Country notes retain all flags.
 Sources: BGS API (4 October 2026); Natural Earth. See the accompanying country table.

Figure 5: Global production in the separate BGS descriptive source

Notes: Source definitions and export-proxy flags remain explicit. Gray is unavailable in the plotted population, not necessarily zero. Each mineral has its own source denominator.

A Contract, procurement and boundary conditions

Let $h > 0$ be the eligible producer quote, $x > 0$ the outside-to-eligible quote ratio, $r \in (0, 1)$ late discovery, $d > 0$ the bond coefficient, $c > 0$ forgery curvature and $\kappa > 0$ diligence curvature. The contract's reporting terms are ℓ_L and ℓ_E . Define

$$A = r\ell_L - \ell_E, \quad a = A/\kappa, \quad (1)$$

$$\pi = \frac{1 - x - rd}{c + d(1 - r)a}, \quad e = a\pi, \quad p = r + (1 - r)e. \quad (2)$$

The reference vector is $(r, d, c, \ell_E, \ell_L, \kappa, \phi) = (.25, .20, 2, .02, 1, .50, .02)$. Supply curvature is .50 and $x_0 = .80$ in the reference comparison. The rest-of-world eligible share is .50; .10 of each focal country's eligible output is vulnerable. These are maintained assumptions. Physical quantities are divided by the world's observed mine-output scale, and monetary accounts by that scale times the source-year reference price. At the baseline, $h_0 = 1$ and $p_{D0} = x_0$.

Let C_{i0}, D_{i0} be each origin's anchored quantities. The new static branches use

$$C_i(h) = C_{i0}[1 - w + wh]^{\eta_{C_i}}, \quad D_i(p_D) = D_{i0}[1 - w + w(p_D/x_0)]^{\eta_{D_i}}. \quad (3)$$

Mining-only gold and single-metal cobalt set $w = 1$; joint-revenue cobalt sets $w = .1727865613$. The fixed recycling origin has $C_{i0} = 1370/3661.2$, curvature zero and no diversion. Safe and outside origins set fraud and contractual buyer cost to zero. At a vulnerable origin the expected buyer-cost coefficient is

$$\mathcal{R}_i = r\ell_L\pi_i - \frac{A^2\pi_i^2}{2\kappa}. \quad (4)$$

The two market equations are

$$E_C = h \sum_i (1 + \phi + \mathcal{R}_i) X_i, \quad E_D = xh \sum_i Y_i. \quad (5)$$

For each supply alternative, both baseline expenditures are computed at $(h, x) = (1, x_0)$; the declared spending multiplier then changes E_D only. Thus the observed physical anchor remains fixed when the quote normalization or recycling closure changes.

The maintained interior branch requires $A > (1 - r)d$, $0 < \pi < 1$, $0 < e < 1$ and physical backing at every vulnerable origin. With permitted output C and candidate output D , entries and outside sales are

$$X = \frac{C}{1 - \pi}, \quad Y = D - \pi X > 0, \quad (6)$$

$$\text{revoked} = p\pi X, \quad \text{live candidate} = (1 - p)\pi X. \quad (7)$$

Thus $X + Y = X(1 - \pi) + D = C + D$. Live entries are $X - p\pi X$, so live entries plus revocations equal X . These identities track material rather than assigning legal status to administrative categories.

For the procurement bridge, impose $K(q) = bq^{1+1/\eta}/(1 + 1/\eta)$ with $b, q, \eta > 0$. The interior supply condition is $P = K'(q) = bq^{1/\eta}$; division by $Pq > 0$ gives

$$\frac{K(q)}{Pq} = \frac{1}{1 + 1/\eta} = \frac{\eta}{1 + \eta}. \quad (8)$$

If the matched procurement share is $\zeta = K(q)/(Pq) + m$, then $(\zeta - m)(1 + \eta) = \eta$, and therefore

$$\eta = \frac{\zeta - m}{1 - \zeta + m}, \quad 0 < \zeta - m < 1. \quad (9)$$

For every positive η there is a compatible $m = \zeta - \eta/(1 + \eta)$ unless independent restrictions exclude it. If $0 \leq m < \zeta < 1$, the conditional set is $0 < \eta \leq \zeta/(1 - \zeta)$. The zero-margin restriction chooses its upper endpoint. The price-stage bridge must hold before applying this algebra to procurement evidence.

Certified expenditure is $(1 + \phi)hX$ plus h times the expected contractual buyer cost. In the evaluated branch that cost is nonnegative and $X \geq C(h)$. Hence $E_C \geq (1 + \phi)hC(h)$ bounds h above by $h_{\max}$. Positive fraud requires $x < 1 - rd$, so the outside quote obeys $xh < p_{\max} = (1 - rd)h_{\max}$. Because $Y \leq D$ and candidate supply increases with its quote,

$$E_D = xhY \leq xhD(xh) < p_{\max}D(p_{\max}). \quad (10)$$

All 302 declared exclusions violate this necessary spending cap. Satisfying it is not asserted sufficient for an equilibrium; a zero-fraud corner is outside this branch.

The zero-output frontier holds the outside quote at its anchored value x_0 (the baseline eligible quote is one). Solve the policy's certified-spending equation with $x = x_0/h$ and recover E_D/E_{D0} .

Candidate output then remains at its anchored value in both power supply and the declared fixed-mix alternative. All 792 scheduled frontiers solve. Their extrema concern the declared discrete design.

B Verification and interpretation

The independent reconstruction classifies all 2,376 declared static cases, solves the 2,074 admissible cases and retains the 302 branch exclusions and checks all 792 frontiers. The primary maximum relative clearing residual is 2.63×10^{-12} ; an independent solver attains 1.59×10^{-14} . Restarts at 56 endpoint cases agree to 9.43×10^{-13} . This is evidence of numerical consistency in the evaluated cases, not a theorem of global uniqueness.

Trust principal obeys $B_t = B_{t-1} + D_t^{\text{new}} - R_t - F_t$, where R_t and F_t are refunds and forfeitures. Independent cohort reconstruction checks these identities, grant timing and the separate service and carrying-charge accounts across 202 market-dates. Corrupted material and account records are rejected. The detailed certificate retains the primitive restrictions and the distinction between conditional output changes and observed policy effects.

C 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.

D 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 4: 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)

E Country extraction sensitivity in the matched-curvature grid

Each row retains the original quote, discovery, diligence and spending grid, with $\eta = 1$ and unknown candidate shares 0.5. These are explored extrema of paired screening effects at common spending, not confidence bounds. The shared outside market can transmit an effect to several origins. “Candidate” remains the assigned stratum, not all artisanal or illegal output.

Table 5: Candidate output effects by country (%)

Mineral	Country	Minimum	Maximum	Pairs
Gold	Peru	-0.01513	-0.00304	32
Gold	Colombia	-0.01513	-0.00304	32
Cobalt	United States	-0.85585	-0.16118	34
Cobalt	Australia	-0.85585	-0.16118	34
Cobalt	Canada	-0.85585	-0.16118	34
Cobalt	China	-0.85585	-0.16118	34
Cobalt	DRC	-0.85585	-0.16118	34
Cobalt	Cuba	-0.85585	-0.16118	34
Cobalt	Indonesia	-0.85585	-0.16118	34
Cobalt	Madagascar	-0.85585	-0.16118	34
Cobalt	Papua New Guinea	-0.85585	-0.16118	34
Cobalt	Philippines	-0.85585	-0.16118	34
Cobalt	Russia	-0.85585	-0.16118	34
Cobalt	Turkey	-0.85585	-0.16118	34
Tantalum	Australia	-0.23014	-0.04290	34
Tantalum	Bolivia	-0.23014	-0.04290	34
Tantalum	Brazil	-0.23014	-0.04290	34
Tantalum	Burundi	-0.23014	-0.04290	34
Tantalum	China	-0.23014	-0.04290	34
Tantalum	DRC	-0.23014	-0.04290	34
Tantalum	Ethiopia	-0.23014	-0.04290	34
Tantalum	Mozambique	-0.23014	-0.04290	34
Tantalum	Nigeria	-0.23014	-0.04290	34
Tantalum	Russia	-0.23014	-0.04290	34
Tantalum	Rwanda	-0.23014	-0.04290	34
Tin	Australia	-0.22952	-0.04278	34
Tin	Bolivia	-0.22952	-0.04278	34
Tin	Brazil	-0.22952	-0.04278	34
Tin	Myanmar	-0.22952	-0.04278	34
Tin	China	-0.22952	-0.04278	34
Tin	DRC	-0.22952	-0.04278	34
Tin	Indonesia	-0.22952	-0.04278	34
Tin	Laos	-0.22952	-0.04278	34
Tin	Malaysia	-0.22952	-0.04278	34
Tin	Nigeria	-0.22952	-0.04278	34
Tin	Peru	-0.22952	-0.04278	34
Tin	Russia	-0.22952	-0.04278	34
Tin	Rwanda	-0.22952	-0.04278	34

Mineral	Country	Minimum	Maximum	Pairs
Tin	Vietnam	-0.22952	-0.04278	34
Tungsten	Australia	-0.22449	-0.04184	34
Tungsten	Austria	-0.22449	-0.04184	34
Tungsten	Bolivia	-0.22449	-0.04184	34
Tungsten	China	-0.22449	-0.04184	34
Tungsten	North Korea	-0.22449	-0.04184	34
Tungsten	Portugal	-0.22449	-0.04184	34
Tungsten	Russia	-0.22449	-0.04184	34
Tungsten	Rwanda	-0.22449	-0.04184	34
Tungsten	Spain	-0.22449	-0.04184	34
Tungsten	Vietnam	-0.22449	-0.04184	34

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

- Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2005. “Institutions as a Fundamental Cause of Long-Run Growth. In *Handbook of Economic Growth*, vol. 1A, 385–472. [Primary source](#).
- Ministerio de Energía y Minas del Perú. 2025. *Anuario Minero 2024*, pp. 66–67. Fine-gold production and administrative strata. [Primary source](#).
- ANM–UPME. 2025. *Boletín Estadístico Informativo Minería en Cifras*, no. 24, December, p. 7. Frozen preliminary 2024 royalty-associated production vintage. [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](#).
- U.S. Geological Survey. 2026. *Mineral Commodity Summaries: Cobalt*, pp. 1–2. Revised 2024 production and refined-price vector. [Primary source](#).
- Gulley, Andrew L. 2023. “China, the Democratic Republic of the Congo, and artisanal cobalt mining from 2000 through 2020. *Proceedings of the National Academy of Sciences* 120(26): e2212037120. [Primary source](#).
- CMOC Group Limited. 2025. *Annual Report 2024*, p. 6 (PDF p. 7). DRC copper and cobalt segment revenue; industrial comparator. [Primary source](#).
- World Gold Council. 2025. *Gold Demand Trends: Full Year 2024*, 5 February. Supply and annual reference-price tables. [Primary source](#).
- OECD. 2025. *Enhancing Regional Mining Ecosystems in the European Union*, Box 3.9, p. 96. Broad mining research-spending benchmark. [Primary source](#).
- U.S. Geological Survey. 2025. *Mineral Commodity Summaries: Cobalt*. Historical CES comparison vintage only. [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.