

Innovation, market segmentation and state capacity

The fiscal limits of formalization

Elies Lucio

Calibration 3 · 4 October 2026

Abstract

This global edition preserves the original innovation and state-capacity model and executes 936 dynamic cases across five mineral markets. It compares supply exponents, pooled and segmented markets, frozen/private/supported research, country intervention and paired policies over model periods 0–100. Observed production anchors are separated from assumed candidate populations, fiscal capture and research technology. Independent full-country equations determine admission, while numerical failures, physical boundaries and unlicensed continuations remain visible. Country cash requirements and central service, research and collateral accounts are kept separate. Alternative nonlinear VCG and Muir versions reconstruct owner incentives and finite continuation on the same evidence bundle. Agreement is conditional on shared primitives, not independent empirical confirmation; policy-effect confidence intervals and p-values remain unavailable.

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 government confronting contested mineral production faces at least three distinct decisions. Buyers must decide which claims to accept. Innovators must decide which production methods to improve. Public authorities must decide whether the future return from intervention compensates for current costs and the rents displaced by formalization. The decisions are connected, but none can be inferred mechanically from another. A cleaner certification channel may coexist with persistent outside production. A research subsidy may lower the future cost of eligible production while leaving the government unwilling to act. A programme that covers its current administrative expenses may still require substantial grants, private collateral and a public liquidity bridge.

The contribution of this calibration is to put those decisions in one internally consistent economy and retain their separate financial accounts. It does not convert administrative mining categories into observations of criminal production. Throughout, **candidate output** denotes the production stratum assigned to the model’s outside technology. Eligibility and candidate status are maintained classifications. They are not direct observations of illegality, armed control, environmental harm or household welfare. This distinction is particularly important because the residual world sector is quantitatively large and its classification is assumed.

The principal substantive result is that the direction of research, the incentive for government intervention and the financing of the transition respond to different margins. Segmentation changes relative research prizes. Effective capture determines how much of the eligible sector’s surplus enters public continuation values. Contractual screening affects admissions and diversion, whereas lower outside expenditure directly compresses the market sustaining candidate production. Their interaction creates useful policy comparisons, provided the comparison is made at common inputs and over an admissible horizon.

2 Measurement and the maintained economy

The gold market is scaled to 3,661.2 tonnes of world mine production and an annual reference price of US\$2,386.2 per troy ounce in 2024. Colombia and Peru participate in this shared market. Their reported production anchors are 56.7 and 108.148345 tonnes. Peru’s candidate share uses reported small and artisanal administrative strata. Colombia’s share carries a 2023 barequero composition into the 2024 scale. The different dates and reporting populations are preserved as measurement restrictions. They are not evidence that a common harmful-production share was observed in both countries. ([World Gold Council, 2025](#); [MINEM, 2025](#); [ANM–UPME, 2025](#); [UNODC and MME, 2025](#))

Cobalt uses the revised 2024 USGS vector: world mine production of 302,000 tonnes of contained cobalt, DRC production of 226,000 tonnes, and a refined-metal reference price of US\$11.84 per pound. The industrial metal quotation values a common accounting unit; it is not an artisanal ore or hydroxide minegate price. The DRC candidate share of 10% is a transported historical ASM midpoint, not a current observation of illegal cobalt. Copper coproduct incentives, ore grades and transaction-specific payability are absent from this dynamic closure. Consequently, cobalt results describe a conditional contained-metal economy rather than an independently validated ASM supply system. ([USGS, 2026](#); [Gulley, 2023](#))

The physical anchors determine initial production coefficients. Demand expenditure is then chosen to reproduce the reference price configuration. The central outside-to-eligible producer quote is 0.80 and the supply elasticity is 0.50. Half of residual world output is assigned to the eligible sector, and one tenth of eligible production in each focal country is vulnerable to claims diversion. These choices determine the market’s exposure to the institution. They are assumptions, not additional observations extracted from national production totals.

Research uses one normalized unit of scientific capacity, with 0.20 reserved for verification and up to 0.80 available to production research. A broad mining R&D/revenue benchmark of 0.5% supplies one spending moment. Conditional on the assumed research success function, laboratory costs, deployment costs, ownership and collateral terms, the model fits innovation-productivity parameters of approximately 0.03773 for gold and 0.03497 for cobalt. The associated reference growth factors are 1.01606 and 1.01488 per model period. These fitted values are not separately identified technology estimates. A spending ratio cannot distinguish productivity from the probability of success, the units of scientific effort or the returns privately appropriated by

innovators. The OECD’s sector benchmark and Rio Tinto’s 2024 company ratio of approximately 0.7417% are therefore used as external spending comparators, with their populations retained. (OECD, 2025; Rio Tinto, 2025)

Time is measured in model periods. The reference year supplies physical and monetary scale; it does not identify the length of a research cycle, the duration of political protection or the calendar interval between intervention attempts. All dates below, including the reporting endpoint of 100, are model dates. Monetary amounts are nominal reference-year USD equivalents, not forecasts of future-dollar budgets.

3 Prices, research and public intervention

In the segmented regime, certified and outside spending clear different sales channels. Diversion moves candidate-origin metal into certified claims, but the model separately records production, admissions, revoked claims and live claims. Screening changes the discovery environment and the incentive to make false claims. It does not mechanically eliminate the underlying output. Producer prices adjust as the two channels clear. Origin backing and eligible-supplier participation must hold; a mathematical price continuation that violates those restrictions is not an economic equilibrium of the maintained institution.

Research owners choose scientific effort against next-period private profits. They do not internalize the effect of their own research on market prices. Eligible innovators also face a collateral wedge: the marginal bond is posted before its discounted refund. Publicly converted producers use previously disclosed eligible technology. The distinction between privately owned improvements and public adoption prevents the same new capacity from being counted twice as an innovation prize.

Private research equalizes active sectors’ marginal returns to science, subject to capacity and nonnegative effort. The supported regime instead implements the full available production-research allocation in the eligible technology using the minimum nonnegative prize consistent with the competing research option. This is a specific implementation instrument, not a generic transfer of a fixed R&D budget. When private incentives already choose that allocation, support has no incremental role and the required prize is zero.

Government intervention is a dynamic choice. An unconverted unit yields protection rents while authorities wait. An attempt incurs a current cost, succeeds with an assumed probability and changes the next-period converted population. A converted unit generates captured tax proceeds net of recurring services. Governments respond competitively to common prices and the other jurisdiction’s path. The relevant comparison is current intervention cost plus foregone protection against the success-weighted continuation advantage. Because that advantage depends on evolving technology, prices and population, a stationary fiscal threshold cannot replace the solved Bellman condition.

The central statutory anchors are 35% for Colombia, 29.5% for Peru and 30% for DRC corporate profit taxation. These are formal-sector rate anchors, not effective mine tax burdens or collection rates. The capture multiplier is separate. The calibration does not infer capture from an EITI reconciliation discrepancy or budget execution rate. Reach, success, protection rents, intervention costs and recurring services remain conditional primitives. Public accounts can supply expenditure scales and declared receipts, but matched liabilities and attempted cohorts are required to identify the corresponding behavioral parameters. (DIAN, 2026; SUNAT, 2026; Democratic Republic of Congo, 2018)

4 The main dynamic comparison

Table 1 holds the central input vector fixed and compares the institutional and research closures with government intervention available. The output figures describe period 100, not growth rates and not standalone policy effects relative to date zero.

Table 1: Period-100 outcomes under the central input vector

Market and closure	Research	Candidate output, tonnes	Eligible output, tonnes	First attempt (period)
Gold, segmented	Private	1,985.0	8,581.5	Colombia 46; Peru 70
Gold, segmented	Supported	1,790.5	9,173.4	Colombia 46; Peru 70
Gold, pooled	Private	6,306.1	3,814.1	Neither by 100
Gold, pooled	Supported	1,974.8	9,067.6	Colombia 71; Peru 93
Cobalt, segmented	Private or supported	56,719.5	1,030,239.5	DRC 66
Cobalt, pooled	Private or supported	63,240.6	1,030,063.1	DRC 92

The gold comparison illustrates the importance of research direction. In the segmented private economy, initial eligible and candidate research allocations are approximately 0.510 and 0.290; candidate research reaches zero at period 45. In the pooled private economy, initial allocations are approximately 0.358 and 0.442, and candidate research remains positive through the reporting endpoint. A common producer price changes the relative private prizes, allowing the candidate technology to develop much further. Supported research changes this allocation from the beginning, producing a substantial difference in accumulated technologies.

This is a comparison of complete maintained regimes. Pooling removes current certification and new bond issuance while inherited cohorts still settle. Its output difference cannot be attributed to a single isolated certification premium. Nor can the greater eligible output be interpreted as a household welfare gain without additional information about costs, labor, income distribution and environmental effects.

Cobalt presents a different margin. Private production research selects the eligible technology from the first date in every explored joint configuration. The supported allocation therefore coincides with the private allocation, and no additional support prize is required. This is an informative redundancy result within the model. It does not establish that real cobalt research is optimally financed, because the assumed ownership, coproduct omission and transported R&D benchmark determine the private returns being compared.

5 Paired policy ranges

The policy experiment compares three states: the baseline, stronger screening, and a combined policy. Screening raises the contractual discovery parameter from 0.25 to 0.50. Combined policy adds a phased reduction in outside expenditure, approaching 25%. The same fitted technology, initial stocks, government parameters and research closure are used on both sides of each contrast. The reported effect is the policy outcome minus its named comparator. A negative production effect denotes contraction.

Four joint vectors were specified before the runs. The central vector uses capture 1, intervention success 0.25, reach 0.25, R&D intensity 0.005 and quote 0.80. A low-capture vector changes capture alone to 0.50. A constrained vector combines capture 0.50, success 0.15, reach 0.10, R&D intensity 0.0025 and quote 0.85. An enabling vector combines capture 1, success 0.40, reach 0.30, intensity 0.007417 and quote 0.70. These are four selected joint configurations, not all points in the rectangle between their coordinates. Composite-vector differences cannot identify a separate causal contribution for each moved primitive.

Table 2 reports explored period-100 extrema relative to the matched baseline. Focal gold output combines Colombia and Peru; focal cobalt output refers to DRC. The distinction matters because the assumed residual world classification generates much of world candidate output.

Table 2: Period-100 paired candidate-output effects

Market and research	Policy	World candidate effect, tonnes	World candidate effect, %	Focal candidate effect, tonnes
Gold, private	Screening	-0.290 to -0.249	-0.01442 to -0.01323	-0.00818 to -0.00497
Gold, private	Combined	-317.338 to -226.087	-14.880 to -12.003	-8.960 to -4.507
Gold, supported	Screening	-0.1205 to -0.1003	-0.00674 to -0.00558	-0.00302 to -0.00226
Gold, supported	Combined	-164.277 to -163.351	-9.13192 to -9.13165	-4.639 to -3.256
Cobalt, private or supported	Screening	-197.773 to -171.264	-0.35341 to -0.28351	-70.342 to -52.920
Cobalt, private or supported	Combined	-5,184.782 to -4,749.260	-8.54049 to -8.48659	-1,933.598 to -1,396.011

Every entry is a **conditional sensitivity range**. None is a confidence interval, identified set or proven global bound. Absolute and percentage endpoints may occur at different vectors because the matched baseline differs across vectors. The complete tables retain endpoint configurations and all planned cases. All four vectors are numerically solved and institutionally admissible through date 101 in these paired comparisons.

Stronger screening contracts candidate production in every explored case, but its scale is small. For focal gold under private research, the range is approximately five to eight kilograms at the reporting date. This is economically distinct from a change in certified claims. The combined policy produces much greater extraction effects because it changes outside spending directly, reduces candidate receipts and alters public intervention incentives. At period 100, candidate receipts fall by approximately 24.98–24.99% for private-research gold and 23.36–23.50% for cobalt. These receipt changes still do not measure lost household income or net social benefit.

The combined policy also brings forward the government’s incentive to act. At the central vector, first attempts move from periods 46 to 28 in Colombia, 70 to 52 in Peru and 66 to 50 in DRC. Screening alone leaves the first-attempt date unchanged in these cases. Halving capture, however, eliminates all attempts through period 100 under baseline, screening and combined policy. The constrained vector does likewise. Thus a favorable market channel does not guarantee that authorities acquire enough of its benefits to intervene. The enabling vector brings first baseline attempts to 13, 25 and 25 respectively, and combined-policy attempts to 5, 14 and 14. These dates reveal decision boundaries among specified economies; they are not calendar implementation forecasts.




Figure 1: Paired dynamic policy effects under the central input vector

Notes: Each line compares a policy with its own matched baseline at the same research closure. Panel scales differ. Private and supported cobalt paths coincide. Periods are not calendar years; lines describe conditional model trajectories without statistical confidence bands.

Statistical significance is **not estimable from current evidence**. There is no matched assignment or sampling design for the policy effect, no estimated joint distribution of the assumed primitives, and no source covariance linking the calibration stages. Neither 101 reporting periods nor four parameter vectors creates an empirical sample. The negative sign is robust across the explored cases; that numerical statement does not supply a p-value or establish robustness over all economically possible parameters.

6 Who pays for the transition?

Operating fees, grants, refundable collateral and public intervention expenses have different payers and timing. Combining them into a single “cost” would obscure the policy problem. The central operating account receives certification fees and pays verification services, research and deployment expenses. Supported research adds success-contingent grants, committed at one date and represented as expected cash at the next. The trust separately receives new collateral, refunds maturing balances and transfers forfeitures. Outstanding principal is a liability-backed stock, not an irreversible programme expenditure. The forfeiture recipient is unspecified, so forfeiture is not silently credited as central revenue.

Country governments receive captured tax proceeds and pay recurring services and intervention attempts. The protection flow represents the incumbent’s modeled rent, not government tax revenue or a measured social resource cost. The private programme-liquidity account records research payments, fees and new deposits, less refunds and grants; it excludes production earnings

and is therefore not firm profit. Scientific wages and laboratory expenses are also distinguished so that compensation transfers are not mistaken for independently measured welfare losses.

In the central private-research cases, fees cover modeled central operating expenditure at each reporting date, so the central cash top-up is zero. This does not make intervention free. The gold government accounts contain an aggregate single-period deficit of about US\$609.5 million at the beginning of intervention; DRC's reaches about US\$158.6 million. Later receipts can produce a positive discounted fiscal balance while an early liquidity requirement remains. Separate country bridges are reported because an arithmetic aggregate cannot authorize transfers between Colombia and Peru.

Research support creates a distinct gold funding need. Under the central baseline, discounted expected grants over periods 0–100 total US\$13.577 billion, whereas discounted positive central cash deficits after fee revenue total US\$3.194 billion. The combined policy reduces those figures to US\$6.877 billion and US\$1.102 billion. The grant total and the fee-net top-up must not be added together: the latter already incorporates the former. These present values use the model's 0.95 per-period discount factor. The calculation is not an annual fiscal projection and does not include an independently estimated cost of bringing about the assumed contraction in outside expenditure.

Collateral introduces another distinction. Net carrying charges of 0%, 2%, 5% and 10% per model period are reported as accounting sensitivities. They are neither observed funding spreads nor alternative equilibria with refitted research incentives. At period 100 in the central segmented private baseline, outstanding principal is about US\$415.3 billion for gold and US\$16.45 billion for cobalt in reference-year value units. If a programme reimbursed a 5% net carrying charge, its fee-net operating-plus-carry requirement at that date would be approximately US\$9.706 billion and US\$397.5 million respectively. The principal itself is not added as another cost. Nor does this reimbursement hypothesis identify which institution should finance private liquidity.

The practical result is a financing frontier rather than a universal recommendation. A policymaker can supply a desired candidate-output contraction and separate limits on central grants, country intervention cash and collateral reimbursement. The reported cases show which explored combinations meet those limits. They cannot establish affordability without actual budgets, borrowing terms and a named payer. They cannot rank welfare without environmental damage, displacement, labor income and alternative-use evidence. In particular, the apparent superiority of combined policy on candidate output does not establish the cost or feasibility of reducing outside demand.

7 What the calibration establishes

The model establishes a coherent conditional mechanism: certification changes private innovation prizes; innovation changes prices and fiscal opportunities; fiscal capture determines whether those opportunities induce intervention. It also establishes that screening, outside-demand restraint and research support are different instruments. Cobalt support is redundant in the explored closure, while gold support materially changes the inherited technology path. Low capture can block intervention even when the production response to the market policy remains favorable.

The independent checks support the executed conditional comparisons. The technical appendices report their numerical scope, the two retained structural boundaries and the separate horizon and state-grid refinements.

The next evidentiary priority is sharply defined. Matched procurement transactions must connect assays, payability, production and research ownership. Attempted intervention cohorts must connect expenditure, success, reversals and production-weighted reach. Tax liabilities and dated collections must identify effective capture on a compatible base. Verification invoices and recovery histories must distinguish operating resources, refundable collateral and collectible penalties. Those measurements would allow the conditional frontier to become an estimated policy comparison. Until then, the complete contribution is a verified calibration with explicit scope, paired ranges and financing obligations—not a causal certificate for deployment.

8 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 3: 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.

9 Five-mineral dynamic extension

The primary design executes 936 cases: thirteen mineral/share contexts, two supply exponents (0.5 and 1), pooled or segmented markets, frozen/private/supported research, government intervention on or off, and three policies. Observed gold strata remain fixed; the other minerals use three assumed unknown-country shares. The baseline, stronger screening and screening with a phased 25% contraction in outside expenditure are paired only at common inputs and a common horizon. Periods 0–100 are model periods without an identified calendar duration.

The larger producer set matters for public decisions. Each positive-production country has its own mass and fiscal account. Colombia’s 0.35, Peru’s 0.295 and DRC’s 0.30 statutory headline assumptions are retained from the benchmark; 0.30 elsewhere is explicitly assumed, not a verified legal or collection rate. Effective collection, intervention reach, success and protection are separate scenario primitives. The ROW residual is not assigned a government.

Research expenditure remains a conditional 0.5% baseline fit. This moment does not identify the innovation-success function, technology improvement or ownership returns. Frozen, privately chosen and supported research therefore answer different hypothetical questions. Government intervention is recomputed from the preserved Bellman problem and actual country path, rather than inferred from a pooled fiscal surplus.

Table 4: Executed primary dynamic design: admission status

Status	Cases
Accepted	495
Equation Verified Predictor	261
Equation Verified Refined	31
Equation Verified Residual	10
Numerical Unresolved	57
Rejected Economic Domain	3
Structural Boundary	78
Unlicensed Continuation	1

Notes: All scheduled cases remain in the register. Equation-verified statuses preserve a false fixed-point-iteration convergence flag where applicable; independent full-country conditions still meet the original thresholds.

Table 5: Period-100 candidate extraction: segmented, no-government bridge

Mineral	Research	Screening (%)	Combined (%)
Gold	frozen	-0.0546	-13.3587
Gold	private	-0.0160	-16.9760
Gold	supported	-0.0089	-13.3841
Cobalt	frozen	-2.3546	-11.4559
Cobalt	private	-0.5604	-12.5850
Cobalt	supported	-0.5604	-12.5850
Tantalum	frozen	-0.8501	-12.8086
Tantalum	private	-0.2805	-17.4615
Tantalum	supported	-0.1520	-13.1774
Tin	frozen	-0.8478	-12.8102
Tin	private	-0.2797	-17.4606
Tin	supported	-0.1516	-13.1780
Tungsten	frozen	-0.8287	-12.8231
Tungsten	private	-0.2727	-17.4536
Tungsten	supported	-0.1479	-13.1833

Notes: Supply exponent $\eta = 1$; unknown shares 0.5. Each entry is a paired difference from its own baseline. A dash means the complete pair is not admitted; it is not zero. The nonlinear rebuild uses its own baseline and finite continuation specification.

10 Country financing, timing and numerical limits

The saved outputs report policy production, eligible and candidate science, first intervention, service costs, fees, grants, deposit vintages, refunds, forfeitures and country tax and protection accounts. The maximum cumulative cash deficit is calculated for each actor separately. Principal held for refund is not fee revenue, and a surplus in one country is not available to another without an explicit transfer institution. Carry sensitivities of 0, 2, 5 and 10% per model period are separate economic accounts.

Some nearly indifferent governments make the original simultaneous best-response iteration cycle. Repeated full-country attempts are preserved. A predictor may be accepted only after the original independent verifier reconstructs the full-country markets, science conditions and Bellman inequalities at unchanged tolerances. Numerical failure, stock/backing failure and an unlicensed continuation are separate statuses. Grid and horizon checks govern the precision with which intervention timing can be discussed; a small on-path residual alone does not prove stability of the first intervention date.

Six of ten selected central segmented-government baselines are admitted at $\eta = .5$ (gold observed shares; other minerals unknown shares .5; private/supported research). Their twelve separate horizon/grid comparisons pass the original equations and stability thresholds: the maximum normalized state difference is 1.2263×10^{-5} and the maximum action difference 1.0279×10^{-4} . Supported tantalum, tin and tungsten, and private tungsten remain unresolved and receive no stability claim. These selected-baseline checks do not establish timing stability for every policy, share, vector or exponent-one path.


10.1 Separate joint institutional sensitivity

A separate 324-case design crosses the original low-capture, constrained and enabling vectors with the three 3T minerals, shares .1/.5/.9, exponents .5/1, private/supported research and three

policies. Markets remain segmented and governments active. There are 209 admitted cases, 99 without licensed continuation and 16 numerically unresolved cases. Low capture changes capture alone to .5. The constrained vector also sets success .15, reach .10, the R&D target .0025 and the initial relative quote .85. Enabling sets capture 1, success .40, reach .30, R&D target .007417 and initial quote .70.

No intervention occurs through period 100 in the 66 admitted low-capture or 84 admitted constrained paths. All 59 admitted enabling paths contain intervention; country-specific first attempts range from 18–48 under baseline/screening and 11–31 under combined policy. Among admitted matched pairs, period 100 combined candidate-output effects range from -17.462 to -7.677% for low capture, -18.069 to -7.711% for constrained and -20.103 to -7.355% for enabling. Admission coverage differs. These finite ranges neither identify the separate effects of jointly moved primitives nor form confidence intervals.


The nonlinear rebuilding changes the source of research rents and the allocation of processing services. It derives the owner prize from its stated prototype rights, installation, deposits and refunds, and funds later service liabilities explicitly. Its accepted finite-horizon cases do not establish the infinite-horizon theorem for all private-research or frozen-technology paths. Differences between the two dynamic packages must therefore be classified by technology, institution, ownership and continuation before being interpreted as a policy disagreement.



Rows: research / government. S = screening; C = combined. Text marks excluded pairs, not zero effects. Period 100; segmented markets; unknown country shares 0.5; observed transported shares unchanged. Finite conditional model comparisons. No confidence intervals or causal estimates.

Figure 2: Period-100 production comparisons: gold, cobalt and tantalum


Notes: Source tables and selection rules accompany the figure. Missing comparisons are not assigned zero.



Rows: research / government. S = screening; C = combined. Text marks excluded pairs, not zero effects. Period 100; segmented markets; unknown country shares 0.5; observed transported shares unchanged. Finite conditional model comparisons. No confidence intervals or causal estimates.

Figure 3: Period-100 production comparisons: tin and tungsten

Notes: Source tables and selection rules accompany the figure. Missing comparisons are not assigned zero.



Combined policy; maximum cumulative cash deficit over model periods 0–100.


Each country dot is a separate account; lines show the country minimum–maximum, not uncertainty.

Axes vary by panel. Zero is a computed zero; text marks missing pairs. No cross-country transfers.

Reference-value normalization compares scales; it is not a share of GDP, annual revenue or welfare.

Figure 4: Separate financing requirements: gold, cobalt and tantalum

Notes: Source tables and selection rules accompany the figure. Missing comparisons are not assigned zero.



Combined policy; maximum cumulative cash deficit over model periods 0–100.
Each country dot is a separate account; lines show the country minimum–maximum, not uncertainty.
Axes vary by panel. Zero is a computed zero; text marks missing pairs. No cross-country transfers.
Reference-value normalization compares scales; it is not a share of GDP, annual revenue or welfare.

Figure 5: Separate financing requirements: tin and tungsten

Notes: Source tables and selection rules accompany the figure. Missing comparisons are not assigned zero.

Evidence coverage: counts, not identification

2024 production and customs records; annual prices cover 2005–2024.

IPIS shows each site’s latest admitted visit through 2024; mineral tags overlap.

	BGS positive countries	Import reporters	Mirror origins	3T route origins	Annual price years	IPIS DRC sites	Core accounts
Gold	102	112	148	—	20	1898	3
Cobalt	19	89	68	—	—	—	13
Tantalum	13	50	70	6	—	203	12
Tin	22	50	57	6	20	674	16
Tungsten	23	38	51	6	—	69	13
Copper	55	84	108	—	20	9	Context
Nickel	28	54	57	—	20	—	Context
Diamonds	24	74	74	—	—	74	Context

Dash: no admitted series, not zero activity. Selected routes query six origins.

BGS includes export proxies; customs origins include territories/residual areas.

IPIS is a selected site panel, not a national prevalence estimate. Core accounts include ROW. Source stages and uncertain parameters remain separate.

Sources: BGS, UN Comtrade, World Bank, IPIS (ODC-BY) and frozen registry.

Figure 6: Evidence coverage across the global resource comparison

Notes: Source tables and selection rules accompany the figure. Missing comparisons are not assigned zero.

A Maintained primitives and units

The values below complete the reference economy described in the text. All are assumed. Policy vectors change only the declared inputs. Research and government discount factors are per model period; costs and rents use the model's normalized value units. The protection parameter ω_P below is distinct from the procurement share ζ in contributions 1 and 2.

Table 6: Assumed central primitives

Block	Values
Contract	$r = .25, d = .20, c = 2, \kappa = .50, \ell_E = .02, \ell_L = 1, \phi = .02$
Supply and exposure	$\eta = .50, x_0 = .80$; vulnerable eligible fraction .10; rest-world eligible share .50
Research	$\delta_R = .95, \nu = \log 2$, science 1, verification .20, laboratory curvatures .002 (both uses)
Deployment and collateral	$K_C = K_D = .001$, owner bond coefficient $b_o = .20$, retention $H = 3$ model periods
Government	$\delta_g = .95$, attempt cost $k = .50$, protection $\omega_P = .40$, recurring service $f = .04$, success $\beta = .25$, reachable fraction $\rho = .25$
Fiscal capture	$\chi = 1$ times the stated statutory anchor; the low-capture vectors set $\chi = .50$

Notes: The cost-growth index is $L_t = g^t$. The parameter ρ in this dynamic model denotes reachable capacity; it is not the Muir institutional parameter in contribution 2. No complete empirical fiscal or cost system is inferred from these values.

B Physical accounts and market clearing

Let eligible and candidate production coefficients be B_s, B_v, B_D , where B_v is eligible production vulnerable to diverted claims. Let A_{Ct}, A_{Dt} denote privately owned technology; public adopters use $A_{C,t-1}$. For country j , the conversion coefficient is

$$w_j = \rho B_{Dj} x_0^\eta, \quad W_t = \sum_j w_j (1 - m_{jt}),$$

where m is the unconverted fraction of reachable capacity, not a share of the entire national population. Eligible and candidate production are

$$Q_{Ct} = (B_s A_{Ct} + B_v + W_t A_{C,t-1}) h_t^\eta, \quad Q_{Dt} = (B_D - W_t/x_0^\eta) A_{Dt} p_{Dt}^\eta.$$

The candidate sector label is a model assignment, not observed illegality. With relative producer quote $x_t = p_{Dt}/h_t$, diversion is $D_t = B_v h_t^\eta \pi_t / (1 - \pi_t)$, so that $X_t = Q_{Ct} + D_t$, $Y_t = Q_{Dt} - D_t$ and $X_t + Y_t = Q_{Ct} + Q_{Dt}$. Revoked claims are discovery times diversion; live claims and live candidate claims must be computed after revocation.

Write $A = r\ell_L - \ell_E$ and $a = A/\kappa$. Under the maintained no-silence condition $A > (1 - r)d$, interior fraud and effort satisfy

$$\pi = \max\{0, (1 - x - rd)/[c + d(1 - r)a]\}, \quad e = \min\{a\pi, 1\}.$$

If effort saturates, fraud is re-solved as $\max\{0, (1 - x - d)/c\}$. Discovery is $r + (1 - r)e$. Expected retained payment is $R^q = \pi r \ell_L - A \pi e + \kappa e^2/2$, distinct from government net return below. Segmented clearing requires

$$E_{Ct} = h_t(1 + \phi)X_t + h_t R_t^q B_v h_t^\eta / (1 - \pi_t), \quad E_{Dt} = p_{Dt} Y_t.$$

Demand is $E_{Ct} = E_{C0} g^t$ and $E_{Dt} = E_{D0}[1 - q(1 - e^{-0.2t})]$, with $q = 0$ except combined policy, where $q = 0.25$. Fixed-expenditure historical sensitivities are separate. Pooled clearing has one producer price and no new certification claims:

$$p_t = \left[\frac{E_{Ct} + E_{Dt}}{B_s A_{Ct} + B_v + W_t A_{C,t-1} + (B_D - W_t/x_0^\eta) A_{Dt}} \right]^{1/(1+\eta)}.$$

Pooled X, Y are origin accounts only. The feasible segmented branch requires $\pi < 1$, country-specific candidate output at least country diversion, and $x \leq 1$ for eligible supplier participation. These restrictions are checked for realized paths and for the hypothetical state grids used by the government continuation. A finite state grid is not a proof over all states.

C Research, spending fit and technology timing

Research success is $z(s) = 1 - e^{-\nu s}$, with $\nu = \log 2$. Available production science is $S = 1 - 0.2 = 0.8$. Technology obeys

$$A_{C,t+1} = A_{Ct}[1 + \gamma z(s_{Ct})], \quad A_{D,t+1} = A_{Dt}[1 + \gamma z(s_{Dt})].$$

Let $v_C = 1 + \gamma z(s_C)$ and let the marginal research collateral be $b_t^I = b_o \gamma I_{C,t+1}^{\text{owned}} / v_C$. The eligible and candidate private prizes are

$$J_{Ct} = \delta_R [\gamma \Pi_{C,t+1}^{\text{owned}} / v_C - K_C A_{Ct} - b_t^I] + \delta_R^{H+1} b_t^I, \\ J_{Dt} = \delta_R [\gamma \Pi_{D,t+1} / v_D - K_D A_{Dt}].$$

Pooling sets marginal certification collateral to zero. The private KKT conditions are $s_k \geq 0$, $w \geq 0$, $s_C + s_D \leq S$, $z'(s_k)J_k - \ell A_k s_k \leq w$, equality for active research, and $w(S - s_C - s_D) = 0$. The supported instrument sets $s_C = S$, $s_D = 0$ and

$$G_t = \max \left\{ 0, \frac{\max(0, z'(0)J_D) + \ell A_C S}{z'(S)} - J_C \right\}.$$

The implementation adds an explicit $10^{-8} A_C$ numerical tie-breaking increment when the prize is positive. This is disclosed, negligible relative to the reported spending scales, and is not an empirical subsidy premium. Expected grant committed at t is $z(s_{Ct})G_t/\delta_R$; it is paid in the expected-flow ledger at $t + 1$. The final reporting-date research conditions therefore require saved market date 101.

The reference spending fit solves a scalar equation conditional on all these primitives: reference research bill equals r_{RD} times reference producer receipts. The reference bill contains wages for normalized science plus the production and verification laboratory terms. It is not a separate observation of γ . The fitted growth index is $g = 1 + \gamma z(0.8)$, used for certified expenditure and

public programme cost scaling. The R&D spending moment, technology parameter and growth assumption must retain separate labels.

D Dynamic government comparison and local frontiers

With public cost index $L_t = g^t$, statutory anchor τ_j , capture χ , recurring-service cost f , and protection parameter ω_P ,

$$R_{jt} = \chi \tau_j \frac{A_{C,t-1} h_t^{1+\eta}}{(1+\eta)L_t} - f, \quad b_t = \omega_P \frac{A_{Dt} p_{Dt}^{1+\eta}}{x_0^\eta (1+\eta)L_t}.$$

Set $d_g = \delta_g g < 1$. Conditional on the other country paths and the technology path, an aggregate attempt rate u implies $m' = m(1 - \beta u)$. The normalized value recursions are

$$U_t = R_t + d_g U_{t+1}(m'),$$

$$V_t = \max\{b_t + d_g V_{t+1}(m'), -k + d_g [\beta U_{t+1}(m') + (1 - \beta)V_{t+1}(m')]\}.$$

The action advantage is $a_t = -k - b_t + d_g \beta \Delta_{t+1}(m')$, where $\Delta = U - V$. At $u = 0$, $a \leq 0$; at $u = 1$, $a \geq 0$; at an interior rate, $a = 0$. This yields the independently reconstructed difference identity

$$\Delta_t = R_t - b_t + d_g \Delta_{t+1}(m') - \max\{0, a_t\}.$$

At a *fixed solved continuation* with positive Δ_{t+1} , the local break-even success probability is $\beta_{req} = (k + b)/(d_g \Delta_{t+1})$, and the maximum current attempt cost is $k_{max} = d_g \beta \Delta_{t+1} - b$. If the continuation difference is nonpositive, nonnegative costs and protection cannot generate a positive action advantage. These formulas do not identify a structural success probability or a global capture threshold. Changing success changes both the state transition and the continuation, requiring a new equilibrium.

`government_break_even_frontiers.csv` reports this local frontier with its solved continuation and explicit limitation. The first-attempt comparison uses rates above 1e-5, the archived reporting convention. The underlying complementarity check retains actual rates and its stated boundary tolerances.

E Cash, principal and economic-accounting identities

Let F_t be fees, C_t operating/verification/deployment expenditure, and G_t^{paid} expected grants paid. Central cash is

$$NC_t = F_t - C_t - G_t^{paid}, \quad A_t = \max\{0, -NC_t\}.$$

Thus grants are already included in A_t ; adding both again double counts. The inherited trust initially holds H outstanding baseline cohorts in both regimes. With deposits D_t^B , refunds R_t^B , forfeiture transfers L_t^B and outstanding principal B_t ,

$$B_t - B_{t-1} = D_t^B - R_t^B - L_t^B.$$

Forfeiture is a trust outflow whose destination is unassigned. It is not central fee income. Private net programme liquidity is deposits minus refunds, plus fees and production research expenditure, minus grants paid. It excludes production earnings and therefore is neither private profit nor welfare cost. Adding forfeitures again to this private net deposit account would count the same retained principal twice.

Country government budget balance is tax received minus recurring services minus intervention attempts. Archived `fiscal_net` excludes the attempt flow, so the revised complete budget subtracts it explicitly. Protection rents are reported separately. Country bridges are not automatically fungible across jurisdictions, and no central-country transfer mechanism is assumed.

The carry sensitivity displays $A_t(c) = \max\{0, C_t + G_t^{paid} + cB_t - F_t\}$, for $c \in \{0, 0.02, 0.05, 0.10\}$. It is a hypothetical reimbursement account if the programme internalizes net collateral carrying charges. It neither spends principal nor measures social resource cost. The rates are per model period and do not re-solve private research incentives. The historical `central_economic_gap` excludes grants; it is retained for compatibility, not used as an all-in programme-cost measure.

The financing table reports both discounted sums of positive period top-ups (no carryforward of surpluses) and the bridge $\max\{0, -\min_t \sum_{v \leq t} NC_v\}$, allowing the same actor to retain undiscounted surpluses at zero interest. These are different financial conventions. Country-specific bridges and their sum are supplied alongside the arithmetic aggregate; the aggregate does not establish intergovernmental transfer authority. Financial flows are deterministic expected paths, not liquidity-risk quantiles from a stochastic programme model.

F Paired effects, explored ranges and inference

For every declared vector θ , policy/comparator pair and research closure, compute $Y_p(\theta) - Y_c(\theta)$. All nonpolicy primitives and fitted nuisance parameters are checked for exact equality within the pair. Each side must remain valid through date 101. Percentage effects divide that paired difference by the same vector’s nonzero comparator; percentages are omitted where the denominator is immaterial, the outcome is a signed government net budget balance, or a percentage-point difference is the defined unit. A separate review found that an absolute-denominator percentage on signed fiscal balances had been insufficiently labelled; those percentages are now omitted entirely, while the verified dollar differences are retained. The range consists of the minimum and maximum of these paired effects over the four declared vectors, not a difference of independently minimized levels. Endpoints for tonnes and percentage effects need not coincide.

Four vectors cannot certify a continuous domain. There is no empirical sampling or assignment basis for a confidence interval or p-value here. Every range row therefore records NOT ESTIMABLE FROM CURRENT EVIDENCE and its missing design, rather than “statistically insignificant.” No arbitrary practical-benefit threshold is imposed. `policy_financing_frontier.csv` permits comparison with a supplied output target and separate actor budget caps.

G Numerical scope, provenance and files

The initial design contains 68 planned cases: 24 factorial baselines and 48 joint policy/vector/research/mineral cases, with four overlaps. Sixty-six were computed and accepted after independent reconstruction. Two frozen segmented BSC cases are explicit structural boundaries: indefinitely growing certified spending with frozen production technology eventually

requires an unspecified contracting continuation. No tail repair or artificial welfare closure was introduced. Historical finite-terminal diagnostics remain historical, not newly executed evidence.

Ordinary fresh runs use $T=500$, $N=301$ and fixed-point tolerance $2e-8$. The enabling vector uses $T=700, N=501$. Frozen segmented/no-BSC runs use $T=110$ because they have no forward government value problem and a distant invalid tail need not be fabricated. Four central segmented innovation baselines separately receive $T=700, N=301$ horizon checks and $T=500, N=601$ grid checks. All eight refinements pass the independent equation checks. Refinement changes are in `refinement_comparison.csv`; solver residuals and refinement differences are different objects.

Static and accounting checks use $1e-8$, dynamic complementarity and on-path Bellman checks $1e-6$. The discrepancy metric is absolute difference divided by $\max(1, \text{absolute actual}, \text{absolute expected})$, so near-zero normalized values receive an absolute check rather than an unstable relative ratio. Output-unit transformations are separately checked. Maximum initial outside-clearing error is $2.78e-13$, research KKT error $2.47e-14$, government complementarity error $1.26e-9$ and Bellman difference error $1.68e-9$. The initial certificate covers 608,202 conditions; the nonoverlapping refinement certificate covers 74,976. The main-only certificate is a subset, not another independent replication.

`provenance.py` rejects old existence-only caching. Each fresh result records hashes of model code, runner, provenance code, the historical imported model, data, experiment driver and predeclaration, together with the scenario, full solver configuration, Python executable/version, NumPy, SciPy and platform. Historical outputs without that certificate are not cache hits. Cache tests change input bytes, scenario, solver settings and runtime identifiers; no protected original input is modified. Independent verifier reports hash its saved-path and metadata inputs. The fresh-versus-historical comparison is labelled as such and is not an independent empirical validation.

H 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.

I 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 7: 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)

J Separate country financing in the central supported experiment

The table reports each actor's maximum cumulative cash deficit through period 100, in millions of nominal 2024 reference-price dollars. It uses segmented markets, $\eta = 1$, supported research, government enabled and unknown candidate shares 0.5. Each policy needs its own admitted baseline pair. A dash means that pair is unavailable, not zero financing. Central service and grant accounts remain separate from the country public accounts. No transfers are assumed.

Table 8: Maximum cash bridge by actor (reference USD million)

Mineral	Actor	Screening	Combined
Gold	Central	2,553.223	701.187
Gold	Peru	2,293.116	1,699.348
Gold	Colombia	1,986.647	1,473.212
Cobalt	Central	0.000	0.000
Cobalt	United States	2.419	1.974
Cobalt	Australia	57.427	47.167
Cobalt	Canada	40.246	33.056
Cobalt	China	24.027	19.735
Cobalt	DRC	543.017	446.015
Cobalt	Cuba	41.448	34.043
Cobalt	Indonesia	420.478	345.366
Cobalt	Madagascar	37.243	30.590
Cobalt	Papua New Guinea	31.596	25.952
Cobalt	Philippines	37.242	30.590
Cobalt	Russia	96.109	78.941
Cobalt	Turkey	26.430	21.709
Tantalum	Central	–	–
Tantalum	Australia	–	–
Tantalum	Bolivia	–	–
Tantalum	Brazil	–	–
Tantalum	Burundi	–	–
Tantalum	China	–	–
Tantalum	DRC	–	–
Tantalum	Ethiopia	–	–
Tantalum	Mozambique	–	–
Tantalum	Nigeria	–	–
Tantalum	Russia	–	–
Tantalum	Rwanda	–	–
Tin	Central	207.092	94.020
Tin	Australia	90.698	73.139
Tin	Bolivia	170.159	137.217
Tin	Brazil	221.528	178.641
Tin	Myanmar	160.528	129.450
Tin	China	569.873	459.549
Tin	DRC	208.686	168.285
Tin	Indonesia	441.451	355.988
Tin	Laos	14.929	12.039
Tin	Malaysia	43.824	35.340
Tin	Nigeria	24.882	20.065
Tin	Peru	1,397.268	1,072.208
Tin	Russia	26.166	21.100
Tin	Rwanda	32.908	26.537
Tin	Vietnam	88.290	71.198

Mineral	Actor	Screening	Combined
Tungsten	Central	60.081	27.126
Tungsten	Australia	7.008	5.628
Tungsten	Austria	6.398	5.139
Tungsten	Bolivia	12.949	10.400
Tungsten	China	510.336	409.889
Tungsten	North Korea	14.472	11.624
Tungsten	Portugal	4.951	3.977
Tungsten	Russia	11.425	9.177
Tungsten	Rwanda	9.902	7.953
Tungsten	Spain	5.332	4.282
Tungsten	Vietnam	25.898	20.800

K 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.

L 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

- World Gold Council. 2025. Gold Demand Trends: Full Year 2024, 5 February. Supply and annual reference-price tables. [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](#).
- 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](#).
- OECD. 2025. Enhancing Regional Mining Ecosystems in the European Union, Box 3.9, p. 96. Broad mining research-spending benchmark. [Primary source](#).
- Rio Tinto. 2025. Annual Report 2024, p. 178, notes 6–7. Group research expenditure and consolidated sales; transported spending comparator. [Primary source](#).
- DIAN. Corporate income-tax guidance, general rate 35%. Accessed 4 October 2026; does not characterize the complete mining fiscal regime. [Primary source](#).
- SUNAT. General income-tax regime guidance, rate 29.5%. Accessed 4 October 2026; a statutory anchor, not an effective mining collection rate. [Primary source](#).
- Democratic Republic of Congo. 2018. Revised Mining Code, Article 247, profits-tax rate 30%. Official legal text mirrored by Leganet. [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.