

When certification changes production: incentives, incidence and state capacity

Elies Lucio

4 October 2026

Abstract

Certification can improve recorded provenance while barely changing extraction. This global edition compares two complementary calibrations for gold, cobalt, tantalum, tin and tungsten: an extension of the established three empirical models and independent nonlinear VCG and Muir reconstructions. A shared 2005–2024 production and trade archive supplies observed scales while unidentified behavioral parameters remain scenarios. The comparison separates raw extraction, certificate composition, physical processed output, household incidence, innovation and actor-specific finance. Matched supply curvature prevents technological changes from being mistaken for mechanism effects. Independent checks verify numerical and accounting consistency, while failed or inadmissible cases remain visible. The global synthesis classifies shared, mechanism-dependent, source-dependent and unresolved conclusions. Related models are not independent empirical confirmations, and causal policy-effect confidence intervals and p-values remain unavailable.

Keywords: mineral certification; institutional economics; policy incidence; innovation; fiscal capacity; calibration.

Related work. Global edition of the journal companion to *For a Handful of Gold*. The original benchmark discussion is retained; the added global synthesis compares the two latest consolidated theoretical parents and the six extended or alternative empirical papers.

Reading this edition

The following benchmark sections preserve the previous gold/cobalt experiments and their original populations. Their numerical claims should not be attributed to the enlarged country registry. The global comparison later in the main text adds tantalum, tin and tungsten, distinguishes both model families and reports the expanded evidence and verification scope. Original papers and results remain unchanged in the preserved replication snapshot.

1 From recorded provenance to production incentives

A certificate changes the terms on which a buyer accepts a mineral. It need not change the return to extracting it: material excluded from one market may still find a buyer elsewhere. This gap between recorded provenance and production is the institutional problem examined here. When does certification alter the marginal return to extraction, who bears the adjustment, and what sustains it? Answering these questions requires following both the material and the payments through the supply chain. It also requires distinguishing the observations that establish the scale of a calibration from the assumptions governing how producers, intermediaries and governments respond.

This paper complements Lucio [2026b], *For a Handful of Gold*, through three calibrations of its mechanisms. The first connects certification contracts and procurement to production. The second follows a restriction through selling arrangements, essential-goods provision and household participation. The third examines how market segmentation affects innovation, government intervention and transition finance. Together they explain why a change in certification can have different consequences for recorded sales, output and income. They remain separate economies, with different instruments, populations and horizons; they are neither a jointly estimated model nor three independent observations of the same policy effect.

The outside market is the first limit on certification's production effect. A sharp fall in candidate material represented by surviving certificates can coexist with a small contraction of extraction if alternative buyers continue to pay for it. When a restriction does reach the producer, selling arrangements and household participation constraints can preserve quantity while intermediary profit falls. Further conditions arise when adjustment depends on innovation or government-led conversion. Research needs an adequate private return or a funded support instrument; public conversion needs captured gains sufficient to make intervention worthwhile. Each activity also needs finance at the date it is undertaken. The low-capture cases show that market adjustment can occur without government attempts, so state action is not a universal prerequisite for a production response.

In the first calibration's heterogeneous mining-only gold economies, stronger certification contracts the selected Peruvian candidate stratum by 0.0056–0.21% and the Colombian stratum by 0.0030–0.055%, with outside spending held fixed within every comparison. These explored extrema depend on transported regional procurement restrictions, including provisional Peruvian evidence and assumed margins; they are not country treatment-effect estimates. In the second calibration, a common-world gold case attributes 90% of displaced certified entries to lower

extraction, yet total candidate output contracts by only 0.016%. The displaced entries amount to just 0.0178% of initial candidate output. This case transports the provisional upper Peruvian procurement ratio to common-world gold under zero recovery. The large incidence share describes the division of a small flow.

The dynamic calibration reaches a similar conclusion over its model horizon. Under private research, stronger screening raises assumed late discovery from 25% to 50% and changes model-period-100 world candidate output by approximately -0.014% to -0.013% for gold and -0.35% to -0.28% for cobalt across four declared joint parameter vectors. Model periods have no identified calendar duration. A combined intervention also assumes a phased reduction in outside expenditure approaching 25%. Its larger output changes therefore cannot be credited to screening alone. Halving effective fiscal capture prevents government attempts within the reporting horizon in all explored policy cases. Research support is redundant for cobalt in the explored closure, because private research already selects the supported allocation. Both are substantive limits on a universal policy prescription.

The calibration contribution is to quantify these mechanisms under stated procurement, institutional and fiscal assumptions. Matched comparisons isolate the changed policy inputs; material accounts separate extraction from certified entries; financial accounts distinguish transfers and refundable liabilities from resource costs. The calculations extend neither the parent's theorems to every institutional setting nor its mechanisms into estimates of an implemented programme. Aggregate production, selected procurement accounts and broad research-spending comparators leave the main response parameters and policy counterfactual unidentified. New causal policy-effect significance remains NOT ESTIMABLE FROM CURRENT EVIDENCE: neither a simulated path nor a set of parameter vectors supplies the missing sampling distribution.

Throughout the paper, *candidate* denotes the exposure stratum assigned to a certification or outside-production account. Administrative small-miner categories, a transported artisanal share and an assumed residual-world split do not measure illegal extraction, armed control or environmental harm. These distinctions also govern the interpretation of outcomes. Lower candidate receipts need not mean lower criminal rents, stable production need not preserve household welfare, and higher eligible production need not reduce environmental damage. Establishing those links would require observations of ownership, livelihoods and harm alongside the material accounts. The current calibration contains no such outcome measures and cannot use certification status as their substitute.

The puzzle is not resolved by assigning a single actor responsibility for the entire supply chain. A buyer can improve its sourcing decision without controlling production at the source. An intermediary can maintain purchases by changing the way it sells or recovers procurement payments. A government can value a larger legal tax base while lacking either the incentive or the cash to acquire it. These choices operate at different dates and on different margins. The relevant unit of analysis is therefore the decision together with its feasible alternative, rather than the institutional name attached to a programme. This perspective also explains why a policy can achieve its direct administrative objective while falling short of a broader production objective. The latter requires an additional, observable transmission mechanism.

The analysis begins with measurement because the meaning of a quantity determines what can be inferred from its change. It then follows the restriction from production returns to the distribution of losses, before examining research, public action and financing. A policy that changes recorded admissions while leaving the outside return nearly intact presents a different appraisal problem from one that changes that return but leaves the government unwilling to act. Positive eventual tax receipts also answer a different question from fees sufficient to cover current operating costs. The sequence keeps these decisions connected without assuming that success at one margin guarantees success at the next.

2 Relationship to the parent and the literature

The companion uses Elies Lucio’s editorial revision of 13 September 2026 of *For a Handful of Gold*. That version develops the regional, selling, certification, aggregation and dynamic mechanisms discussed below. Its source dependencies and matching 134-page PDF are identified in the accompanying version register; the older final-output PDF is a different edition. The distinction matters for the interpretation of the calibration: each selected mechanism carries its own market structure, participation conditions and accounting definitions. The numerical exercises use those specified mechanisms rather than treating every result in the parent as applicable to every alternative economy.

Several parent results organize the calibration. The regional analysis links procurement to household choices and an intermediary’s control over imported necessities. Its selling analysis allows the receipt schedule to respond to regulation through rationing, resale, buyer competition and retention. Its exact aggregation result derives power supply and a separate profit coefficient from regional implementation costs. Its common-price certification equilibrium distinguishes production, diverted admissions, outside deliveries and surviving claims. Finally, its dynamic analysis compares private research prizes with governments’ incentives to exchange protection income for captured fiscal returns. These are inherited theoretical mechanisms, not propositions newly proved by the numerical exercises.

Calibration adds observable scales and explicit parameter restrictions to selected implementations of that theory. The production exercise introduces regional procurement margins, heterogeneous candidate supply, separately cleared recycling, joint-revenue alternatives and matched spending changes. The transmission exercise compares each institution with its own baseline, derives an incidence account and decomposes recorded trade shares. The dynamic exercise supplies paired policy paths, research-support and capture comparisons, and financing accounts for each actor. These additions establish the consequences of the specified finite experiments. They do not establish the parent’s existence, uniqueness or long-run results under alternative technologies that fall outside the relevant theoretical assumptions.

The emphasis on the allocation of returns follows a central concern of institutional economics. Acemoglu et al. [2005] argue that institutions shape production incentives and distribution, while distributive consequences help determine institutional choice. Here those relationships take the form of buyer payment rules, household outside options, collectible liabilities and government

claims on eligible surplus. Which actor receives a marginal gain or bears a loss matters for the response of production, research and intervention. An institutional label alone cannot determine that response: the payment rule and the feasible alternative must be specified at the level of the decision.

The selling and innovation mechanisms also have identifiable antecedents. Loertscher and Muir [2022] analyze how nonconcave revenue, rationing and resale affect a seller’s choices. The maintained Muir illustration below belongs to a specified separated-buyer environment; it is not a complete implementation of every mechanism in that literature. Acemoglu [2002] distinguishes price and market-size forces in directed technical change. The dynamic mineral exercise uses private prizes to allocate research across technologies, while retaining ownership, scientific capacity and collateral assumptions. Neither reference supplies the empirical elasticity or research productivity required to turn the calibration into an estimate.

Fiscal incentives depend on more than the legal announcement of a policy. Besley and Persson [2009] model investments in legal and fiscal capacity as politically conditioned choices. Here effective capture and intervention success enter a government’s continuation problem, rather than serving as observed national indices of state capacity. A larger potentially taxable surplus may be difficult to collect, and collecting it may still be less attractive than preserving current protection income. Separating these possibilities permits a comparison of public incentives with the private market response. It also makes clear which additional evidence would be needed to attach a country’s institutional interpretation to a parameter.

Evidence about conflict underscores the need to measure final outcomes independently. Parker and Vadheim [2017] study early US conflict-mineral legislation and related restrictions using eastern-DRC mining and conflict observations from 2004–2012. Their findings motivate attention to displacement and unintended consequences. Their estimand, minerals, locations and period are not the cobalt counterfactual studied here, and their estimates or significance levels are not transferred into this paper. The companion’s contribution is consequently a conditional account of mechanisms and financial feasibility, not a synthesis of causal estimates from unrelated settings.

3 Measurement, model boundaries and matched comparisons

Physical production and price records establish the scale of the exercises. Peru’s 2024 recorded fine-gold production is about 108 tonnes, with selected small and artisanal administrative categories totaling 20.3 tonnes. Colombia’s production anchor is 56.7 tonnes in the frozen preliminary royalty-associated vintage. A share calculated from a separate 2023 barequero table is transported to that anchor, giving 30.5 scenario tonnes. World gold mine production is about 3,660 tonnes. Each number retains its reporting definition; they do not combine into an observed common illegal-production measure [Ministerio de Energía y Minas del Perú, 2025, ANM–UPME, 2025, UNODC and Ministerio de Minas y Energía, 2025, World Gold Council, 2025].

Cobalt uses the revised 2024 vector of 226,000 tonnes for the Democratic Republic of the Congo and 302,000 tonnes worldwide. A historical artisanal scale is transported as a 10% DRC

candidate share, giving 22,600 scenario tonnes. These are contained-cobalt units. A refined-metal reference quotation does not measure the minegate price of ore, hydroxide or an assayed lot. Copper credits, grade, moisture, payability, transport and finance may all matter for that bridge. The fixed-mix alternative in the static exercise examines one explicit coproduct closure; the dynamic cobalt model does not acquire that closure merely because both exercises concern the same metal [U.S. Geological Survey, 2026, Gulley, 2023].

The country coverage is thus asymmetric by design. Peru and Colombia interact in a common gold market, with a residual world sector. DRC cobalt belongs to a separate cobalt market. Half of residual world production is assigned to the eligible sector, and one tenth of eligible production in each focal country is vulnerable to false admissions. Those shares are chosen. Their role is consequential: a world candidate quantity contains a large assumed residual component. A common-market effect cannot be relabeled a national effect, and a national administrative denominator cannot silently replace the world denominator used in a reported comparison.

Baseline quantities identify scale only conditional on the response technology. For any positive reference price and quantity and any positive power-supply elasticity, a coefficient can be chosen to reproduce that observation. The fit therefore does not estimate the elasticity. Likewise, demand expenditure is inverted to reproduce the assumed initial relative quote. Changing that quote without refitting the baseline would alter physical production before any policy is introduced. Contribution 1 explicitly refits expenditure at each declared initial quote, preserving the admitted physical anchor. This ensures that its contrasts concern changed contracts or spending, rather than an unnoticed change in normalization.

Procurement observations constrain a relation among unknowns. A Colombian regional account reports a purchase-to-benchmark ratio of 0.70. The Peruvian 0.85–0.90 interval remains provisional: the accessible passage is indexed evidence and the nominal archived PDF is HTML; underlying transaction dates and sampling details are unresolved. These sources describe selected price levels, not supply responses. Even if assay, deductions and population could be matched, a cost model is still required to infer response curvature [UNODC and Ministerio de Minas y Energía, 2025, Banco Central de Reserva del Perú, 2026].

Under the maintained power-cost envelope, production cost as a share of receipts equals $\eta/(1 + \eta)$. If the procurement ratio ζ includes a net margin or recovery share m , the bridge is

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

Zero recovery selects an endpoint; it is not an observation. Nonnegative recovery permits elasticities approaching zero and gives only a conditional upper endpoint. If negative recovery is admissible, that upper restriction need not hold. Contribution 1 applies the transported curvature to the named candidate origins while leaving other supply at its declared curvature. Contribution 2 retains a separate common-curvature screening illustration. The difference is material and prevents their numerically similar incidence and procurement ratios from being treated as independent validation.

Innovation and government parameters have analogous limits. A broad mining research-to-revenue ratio supplies one scalar spending moment. Conditional on the success function,

scientific units, laboratory expenses, ownership and deployment costs, it can fit one productivity parameter. It does not jointly identify all those ingredients or the calendar duration of a model period. The dynamic exercise uses a 0.5% benchmark and a company-level comparator in sensitivity analysis, preserving their different populations [OECD, 2025, Rio Tinto, 2025]. A statutory tax-rate anchor similarly does not observe effective collection. Capture is separately assumed, while the wider fiscal regime, liability base and collection timing require measurement.

Five categories distinguish the inputs. Observed quantities retain the definitions of their sources. Calibrated quantities solve an accounting or equilibrium restriction conditional on other inputs. Transported quantities apply a relation from another population or date. Assumptions describe behavior or institutions for which the necessary observations are absent, while chosen policy quantities define the intervention. These categories concern what the evidence supports, not how many decimal places a value carries. A precisely stored assumption is still an assumption. The accompanying registers preserve the full values and their classifications; the text rounds them to the precision needed for interpretation.

The same care applies to the meaning of a range. An explored range reports the least and greatest effect among the specified admissible cases. It does not describe the frequency of intermediate effects, a probability that the sign will change, or the chance that a country will attain an endpoint. A parameter-sensitivity range may be informative about a maintained mechanism even when the probability distribution of its inputs is unknown. The appropriate use is to identify which unresolved primitive matters for a decision and what observation could restrict it. Choosing more grid points without a source-based restriction would improve numerical coverage of that design, but would not itself improve empirical identification.

Reproducible arithmetic does not resolve a measurement mismatch. Valuing administrative production at a refined-metal benchmark gives a consistent nominal scale, but it does not reveal cash retained at the mine. Applying an earlier category share defines a scenario, but does not establish that the share persisted. An industrial revenue weight specifies a joint-revenue comparison, but does not establish an artisanal production technology. Each calculation remains useful within that scope. Stronger interpretations require data on the missing relationship—minegate deductions, changes in category composition or the value of coproducts—rather than additional digits in the computed result.

All effect statements use a matched comparator. The static exercise separates a contract change at common outside spending, a spending change under the reference contract and their combined contrast. The village exercise compares each institutional configuration with its own no-charge optimum. The dynamic exercise compares policy and baseline at identical nonpolicy primitives and research closure, over a common admissible reporting horizon. Absolute effects add along appropriate paths; percentages with different denominators do not. The paper never averages across these models, splices extrema from unmatched economies or ranks the interventions by comparing unlike horizons and populations.

4 From credible claims to the return to extraction

The first mechanism concerns a seller’s opportunity after certification changes. The contract has an eligible producer quote and a separate outside quote. Vulnerable origins can divert candidate material into eligible accounts, while buyers choose costly diligence and late discovery can revoke claims. Fraud is an equilibrium response to the relative quotes, contractual liabilities and discovery environment. The physical constraint is local: a vulnerable origin must actually possess the candidate material backing its diverted admissions. Aggregate clearing cannot repair an origin-level shortage.

At a vulnerable origin with permitted production C , candidate production D and fraud fraction π , certificate entries are $X = C/(1 - \pi)$ and outside sales are $Y = D - \pi X$. Production obeys $C + D = X + Y$. If discovery probability is p , surviving candidate claims are $(1 - p)\pi X$; revoked claims are $p\pi X$. A more credible contract can lower surviving candidate claims through lower fraud or greater discovery while the outside market continues to absorb material. There is no accounting identity equating a revoked claim with an unproduced unit. Appendix A specifies the contract, material constraints and both market-clearing conditions.

The reference contract assumes a late-discovery probability of 25%. The static policy grid increases it to 35% or 50% and either retains or reduces diligence curvature. Initial outside-to-eligible quotes are 0.70, 0.80 and 0.85. Outside expenditure is multiplied by 0.80, 1 or 1.20 in separately identified stresses. The contract-only comparison holds spending constant within each pair. Its purpose is to isolate the modeled contractual margin, not to assert that demand would be fixed after a real policy. The combined comparison reveals how much that unresolved demand response matters.

Table 1: Candidate production under contract and spending changes

Candidate stratum and supply closure	Contract only (%)	Combined (%)
Peru gold, heterogeneous mining	−0.212 to −0.00561	−77.82 to 205.6
Colombia gold, heterogeneous mining	−0.0549 to −0.00301	−32.56 to 36.05
DRC cobalt, single metal	−0.743 to −0.0573	−10.25 to 8.68
DRC cobalt, joint revenue	−0.191 to −0.0129	−2.87 to 2.49

Explored extrema over admissible cases, relative to each matched comparator. Contract only keeps outside spending unchanged within the pair. Combined also uses a chosen outside-spending multiplier of 0.80, 1 or 1.20. Gold uses regional procurement-curvature transport, including provisional Peruvian evidence and assumed recovery. Cobalt joint revenue transports an industrial revenue weight. These are candidate exposure strata, not measured illegal production. Source: revised C1 headline range table; exact crosswalk in Appendix D.

Table 1 shows small contract-only extraction changes and sign-changing combined contrasts. The unusually large positive Peruvian endpoint is a consequence of high transported curvature applied to a relatively small stratum within a shared market. It is not a forecast of a three-fold response by Peruvian miners. The negative and positive endpoints together expose the dependence on outside purchasing power and the procurement bridge. Removing the expansion because it appears implausible would conceal precisely the empirical weakness the exercise is meant to display.

Two alternative supply accounts qualify the result. Gold is also cleared with 1,370 tonnes of fixed permitted recycling, added with its baseline demand backing. The fixed flow does not

estimate a recycling elasticity or the release of above-ground stocks. Cobalt joint revenue uses a cobalt share of approximately 17.3% in an industrial copper-cobalt revenue comparator and fixes the other-metal price and physical mix. Supply responds to the resulting revenue index, so a change in the cobalt quote has a smaller local effect on joint revenue. This explains the more muted cobalt range under that closure; it does not establish the actual artisanal coproduct technology [World Gold Council, 2025, CMOC Group Limited, 2025].

Of 2,376 scheduled static cases, 2,074 are valid and 302 violate a necessary outside-spending bound for the maintained strict positive-fraud branch. No scheduled admissible case remains unsolved. The rejected cases are not observations with zero effect and do not disappear from the design denominator. A zero-fraud corner would require its own equilibrium specification. The 792 separately evaluated zero-output frontiers ask how outside spending could restore initial candidate production under the altered contract. They are conditional frontiers in the declared design, not globally certified bounds.

These production results turn the question toward transmission. A certified sale affects extraction through its contribution to marginal receipts and the seller’s remaining opportunities. The live certificate share measures the quality of the register; lost entries measure displacement from it. Neither quantity, by itself, measures the extraction response. A production evaluation must therefore link entries to alternative sales, prices and physical output for the same origins and dates. This requirement becomes more demanding where an intermediary controls both the sale of mineral and the provision of necessities, since a changed receipt can be absorbed elsewhere in the transaction.

5 Who bears a restriction when production is resilient?

The transmission problem first appears in a simple physical identity. Write candidate output as $Q = O + E$, where O is outside sales and E is candidate entries into certification. For a comparison that reduces entries, define $T = (Q_0 - Q_1)/(E_0 - E_1)$. Its denominator is displaced candidate entries, not total output, household income or all certified sales. With fixed outside expenditure and power supply, the exact finite expression in Appendix B separates lower extraction from additional outside sales. The local incidence is $\eta Q/(\eta Q + O)$, which depends on both supply curvature and the initial outside-sales share.

Table 2: Screening incidence and aggregate production

Common-market screening case	Output change (%)	Lost entries / Q_0 (%)	Incidence (%)
Gold reference	−0.005913	0.01774	33.34
Gold, upper Peru procurement transport	−0.01600	0.01778	90.00
Cobalt reference	−0.2860	0.8455	33.83

Assumed late discovery increases from 25% to 50% with other joint inputs unchanged. The Peru-named row transports the regional procurement ratio to common-world gold curvature under zero recovery; it is not a national Peruvian effect. Incidence is lost candidate output divided by lost candidate certified entries. The output denominator includes assumed residual-world candidate assignments. Source: C2 screening transmission records.

The arithmetic in Table 2 is $(Q_0 - Q_1)/Q_0 = T(E_0 - E_1)/Q_0$. Incidence of about 90% on

a displaced flow of 0.0178% of output therefore implies a contraction of about 0.016%. The resemblance between the incidence share and the procurement ratio is partly imposed by the procurement-to-curvature relation. Allowing nonnegative recovery lowers admissible curvature and local incidence. This common-world illustration consequently provides neither a large aggregate contraction nor an independently estimated Peruvian transmission coefficient. Its purpose is to show why the affected flow and the share of that flow borne by production must be reported together.

The separate village exercise asks what happens between downstream receipts and household supply. Households choose mining labor, necessities and retained cash subject to a budget and an outside participation option. An intermediary purchases mineral and, where permitted, chooses a margin on essential goods. The consolidated implementation cost equals the goods' resource cost plus retained household cash less outside wealth, together with transport. Procurement expenditure and recovered goods margins therefore cannot be counted as separate resource costs. Participation constrains the intermediary's output choice throughout this calculation.

Three switches generate eight configurations: selling arrangements, the integrated goods interface and subsistence dependence. Removing subsistence dependence supplies basic consumption independently; it changes outside resources and is not a welfare-neutral relabeling. Each configuration is compared with its own no-charge optimum. Otherwise, a reported policy difference would confound the restriction with the institution's effect on the initial allocation. Appendix B gives the utility function, budget, cash branches, procurement bounds, outside option and complete finite selling economies, including exit.

In the maintained Muir case, the optimized selling rule makes a covered-buyer charge subtract a constant from receipts over the supported quantity branch. Its derivative with respect to output is unchanged. Fixed-route sales instead expose a constant fraction of every additional unit, so the charge reduces marginal receipts. Continued operation and the global comparison with other feasible branches are necessary for output invariance. In the reference example, the empty configuration contracts by 8.5×10^{-5} from about 0.130 normalized units under a charge of 0.10, while the full configuration remains at about 0.319. The full configuration avoids this small benchmark contraction; the invariance applies only on the stated branch.

The finite VCG economy uses another payment rule and another buyer population. Buyers pay the externality they impose on others; removing a whole buyer includes its losing offers. A charge on a winning buyer may leave the rival bids determining payment unchanged, whereas covering a losing buyer can change those externalities. The reference empty configuration contracts by 8.0×10^{-4} from about 0.170 normalized units, while the full configuration remains at about 0.325. These levels cannot rank Muir against VCG because the buyer economies differ. They demonstrate how the same form of charge can enter distinct marginal-receipt schedules.

Output resilience does not imply an absent burden. Intermediary profit can fall while a participation boundary holds the least-cost active allocation in place. In some sensitivity configurations, output is zero both before and after the charge; zero change then describes inactivity, not preservation of an operating mine. The decomposition of absorbed contraction also admits negative channel contributions. Covering a losing VCG buyer can make the

selling contribution negative, and the interface ranking is sensitive to finite package resolution. Keeping these results is necessary to reject a universal ranking without discarding the narrower mechanisms.

Recorded trade destinations provide a useful warning about exposure measurement. In the frozen HS810520 cobalt records, a symmetric decomposition over six comparable African origins assigns +2.09 percentage points to within-origin routing changes and -7.24 points to changing origin weights for the combined EU27, US, Canada and Australia share. The comparable-cohort change is -5.15 points; a separate -0.0587 -point bridge connects it to the full available-origin aggregate. An aggregate decline therefore does not mean every origin moved in the same direction. These are recorded import-value shares, not contained cobalt, final consumption or effective legal coverage [United Nations Statistics Division, 2026].

Whether output resilience is desirable depends on the policy objective and the distribution of its burden. Stable aggregate output does not establish that household consumption has been protected. Conversely, a mechanism that preserves production may frustrate an attempt to reduce a harmful activity even if it lowers an intermediary's rents. Evaluating either possibility requires welfare and distributional weights that the calibration does not supply. The institutional comparisons instead identify the transactions through which a restriction is absorbed. In particular, necessities supplied independently to the household have a resource cost and a payer; their contribution to participation cannot be treated as costless household protection.

The connection to production is an empirical task. Buyer obligations must be linked to actual offers, transactions and outside options; losing and no-bid offers matter for the payment mechanism. Household evidence must establish whether a restriction changes consumption, cash, labor or participation. Destination geography alone supplies none of these links. The transmission results consequently inform which accounts to observe, while the dynamic analysis asks whether a persistent change in those accounts can support technological and public adjustment.

6 Innovation and the public incentive to intervene

The third calibration connects certification to continuation through prices. In segmented markets, eligible and outside spending clear separate channels. Private owners direct research toward the technology offering the higher marginal prize, subject to scientific capacity. Publicly converted production adopts previously disclosed eligible technology. Governments compare intervention with waiting, taking the commodity price path and the other jurisdictions' decisions into account. Certification, technology and state action interact within this dynamic economy, but its research and fiscal assumptions are not imported into the static village model.

The initial production coefficients use the physical anchors already described. Research has one normalized unit of scientific capacity, with 0.20 assigned to verification and 0.80 available to production research. Success is a concave function of effort. Eligible owners value next-period privately owned profits net of deployment expenses and the difference between the cost of posting collateral and its discounted refund. Candidate owners face a separate prize. Public adoption is excluded from the private eligible prize, preventing converted capacity from being counted

again as an invention owned by the researcher. The essential first-order and complementarity conditions are stated in Appendix C.

Supported research is a precisely defined instrument: the minimum nonnegative prize needed to allocate all available production science to the eligible technology. It is not an unspecified research budget or a universal policy recommendation. In the explored cobalt economies, private research already selects that allocation, so the support prize is zero and supported and private paths coincide. For gold, the competing candidate-research opportunity makes support consequential. This contrast is economically informative precisely because it limits the case for a common subsidy across commodities. Ownership and the omission of dynamic copper joint production remain maintained restrictions.

A government that waits receives protection income from the unconverted activity. An attempt has a current cost, succeeds with a specified probability and changes the next period’s converted population. Converted activity produces captured fiscal receipts net of recurring services. The decision therefore compares current attempt cost plus foregone protection with the success-weighted continuation advantage. Effective capture is separate from the statutory anchor. The calibrated statutory rates are 35% for Colombia, 29.5% for Peru and 30% for DRC, used only as frozen formal-sector anchors rather than complete effective mining tax systems [DIAN, 2026, SUNAT, 2026, Democratic Republic of Congo, 2018].

The reference experiment retains baseline, stronger-screening and combined-policy states. Stronger screening raises late discovery from 25% to 50%. Combined policy additionally imposes an outside-expenditure multiplier $1 - 0.25(1 - e^{-0.2t})$. It approaches a 25% reduction; neither the means nor the resource cost of achieving that spending path is estimated. Four joint vectors vary capture, intervention success, reachable production, research intensity and the initial quote. Policy and comparator use the same vector, fitted productivity and research closure. The extrema across those vectors are explored scenario extrema, not confidence intervals or bounds on all intermediate values.

Table 3: World candidate output at model period 100

Market	Research	Screening (%)	Combined (%)
Gold	Private	−0.01442 to −0.01323	−14.88 to −12.00
Gold	Supported	−0.00674 to −0.00558	$\simeq -9.13$
Cobalt	Private or supported	−0.3534 to −0.2835	−8.540 to −8.487

World candidate-output percentage differences at model period 100, relative to each vector’s matched baseline. Combined includes stronger screening and an assumed phased outside-expenditure contraction approaching 25%. Periods have no identified calendar-year conversion. Cobalt support is redundant in the explored closure. Both supported-gold combined extrema round to -9.13% ; the unrounded range remains in the source table. Four joint vectors; no empirical significance or confidence coverage. Source: C3 paired effect ranges.

Screening alone produces much smaller output changes than combined policy in Table 3. The difference arises because outside expenditure directly supports candidate production; changing it also changes the incentives for research and public intervention. The combined contrast therefore cannot be attributed to stricter certification. Its unmeasured implementation cost prevents a welfare ranking against screening alone. Geographic scope matters as well: these are world effects, and their denominator includes the assumed residual-world candidate sector.

Focal-country comparisons must be taken from their own matched records, rather than inferred by assigning a world percentage to national output.

Capture imposes an additional condition. At the central vector, combined policy moves the first modeled attempts from periods 46 to 28 in Colombia, 70 to 52 in Peru and 66 to 50 in DRC. Screening alone leaves these first-attempt dates unchanged. Halving capture eliminates attempts through period 100 under all three policy states, and the constrained joint vector does likewise. A favorable market response therefore does not ensure that enough value enters the government's continuation account to induce action. These are model decision dates, not forecasts of implementation years or estimates of an administrative delay.

The fiscal comparison also explains why a favorable market movement cannot be read as an automatic expansion of state capacity. Higher eligible returns enter the government's objective only through the captured share, and an attempted intervention changes future states only through the assumed probability of success and reachable population. Increasing one component while another remains weak may leave the action choice unchanged. The joint-vector results are useful because they demonstrate such conditional failures without attributing a separate empirical effect to each component. They are not an instruction to increase a statutory tax rate: altering a rate could change behavior, liabilities or collection in ways outside the maintained comparison.

The reference year fixes physical quantities and monetary units. It does not establish how long research takes, when an intervention succeeds or how collateral maturity maps into calendar time. Period 100 is a common reporting point within the model, not a century or a measured implementation deadline. The same distinction applies to discounting: a factor per model period cannot be used as an annual rate without a time bridge. Dated research and intervention cohorts would make the timing assumptions testable. Until those data exist, the paths distinguish private adjustment from public incentives but provide no calendar forecast of either.

7 Financing the adjustment is a separate condition

The same return that induces a decision need not pay its bill when the bill falls due. Certification requires operating services, innovators may require grants, participants may post refundable collateral, and governments may incur intervention expenses before receiving fiscal returns. A programme can satisfy one of these accounts and fail another. The dynamic calibration therefore records central operations, grant payments, the collateral trust, private programme liquidity and country budgets separately. It reports expected-flow paths rather than liquidity-risk quantiles from an estimated stochastic programme.

The distinction begins with grants and cash deficits. Central net cash equals fees less operating, verification and deployment expenses less grants paid. A positive top-up is the negative part of that balance. Grants are consequently already included in the top-up. Adding the gross grant total to the fee-net deficit would count the same obligation twice. Conversely, a zero central deficit means only that fees cover this defined central account at the relevant dates. It does not finance country intervention, compensate private collateral holdings or establish that any actor

can borrow against future receipts.

Table 4: Grants and central financing requirements

Central supported-gold account	Baseline	Combined policy
Discounted expected grants (US\$ billion)	13.58	6.877
Discounted positive cash deficits after fees (US\$ billion)	3.194	1.102

Central vector, periods 0–100, discount factor 0.95 per model period, nominal reference-year USD equivalents. The fee-net deficits already incorporate grants: rows must not be added. Combined policy assumes the phased outside-expenditure contraction; its implementation cost is unmeasured. These are expected model accounts, not calendar budgets or an affordability certificate. Source: C3 actor-finance records.

Table 4 illustrates the difference in the central supported-gold case. Baseline expected grants have a discounted total of US\$13.6 billion, while discounted positive fee-net deficits total US\$3.19 billion. Under combined policy these become US\$6.88 billion and US\$1.10 billion. The model’s per-period discount factor is 0.95. The monetary scale makes obligations comparable within the calibrated accounting system; it does not make them approved national appropriations. In particular, the lower combined-policy financing figures exclude an independently costed mechanism for reducing outside expenditure.

Refundable collateral is different again. Deposits create an outstanding stock with a liability; maturing cohorts split into refunds and forfeitures. Principal is not service income or an additional irreversible resource flow. Forfeiture is a trust outflow with an unspecified recipient and cannot be silently credited to the central fee account. A participant’s net programme liquidity includes deposits less refunds, fees and research payments less grants, but excludes production earnings. It therefore measures neither profit nor household disposable income. These distinctions prevent apparent financial surpluses from being created by inconsistent consolidation.

The carry sensitivity asks what would happen to a reimbursement account if a programme internalized assumed net charges on outstanding collateral. Rates of 0%, 2%, 5% and 10% are per model period, not observed annual spreads. They do not re-solve private research choices under new funding terms. Reporting such an account is useful for identifying a potential payer and obligation, but it does not establish the opportunity cost of the underlying resources. Transfers between participants, the trust, service providers and governments need a consolidated resource interpretation before they can support a social-cost claim.

Country budgets preserve geography and timing. Captured taxes, recurring services and intervention attempts determine each government’s cash path; protection rents remain separate. A later positive discounted balance can coexist with an earlier deficit. The sum of positive period top-ups assumes no carryforward of surpluses, whereas a cumulative cash bridge allows an actor to retain its own undiscounted surpluses at zero interest. Neither convention automatically permits Colombia’s receipts to finance Peru’s obligations or central fees to finance DRC intervention. An actual financing plan requires specified transfer authority, lenders, terms and risk allocation.

These accounts also separate a financing shortfall from an incentive shortfall. Providing an earlier loan could relax a cash constraint without changing a government’s valuation of formalization, whereas changing capture could alter the desired action without supplying the needed cash. The current model evaluates continuation incentives and reports financial

paths; it does not solve a lender’s screening, default or collateral-enforcement problem. A financing instrument cannot therefore be inserted into the reported frontier as if its terms were already estimated. The frontier tells the researcher which actor and dates to investigate and which obligations to avoid counting twice. An operational financing design would be a further institutional specification with its own evidence requirements.

Appraisal should match each proposed activity to its incentive and financing requirements. A production claim needs transmission to the return supporting extraction. Innovation-dependent adjustment needs a research response, and government-led conversion needs an incentive to intervene. Low-capture cases nevertheless retain market output effects without government attempts. A financing plan must specify who pays and when for the activities actually proposed. Given an output objective and separate budget constraints, the financing frontier can identify which obligations require further investigation. It does not choose the objective or establish affordability. Nor can output losses, intermediary profits, household constraints, grants and liquidity needs be combined into a welfare measure without an account of the underlying resources and transfers.

8 Robustness, negative findings and empirical limits

A conditional calibration is most informative when its boundaries are visible. The static design retains all excluded cases and does not substitute an unmodeled corner. The village comparisons preserve inactivity, negative channel contributions and sensitivity to package resolution. The dynamic design retains two structural boundaries: indefinitely growing certified spending with frozen production technology eventually requires an unspecified segmented contracting continuation. No artificial tail is added to certify those cases. Accepted dynamic policy pairs remain valid through the additional market date needed to evaluate final reporting-date research choices.

Numerical verification is conducted separately for each economy. An independent implementation reconstructs static case classifications, admissible equilibria, matched differences and frontiers. The institutional checks reconstruct household choices, finite allocations, whole-buyer VCG payments and channel decompositions. The dynamic checks cover saved-state identities, model equations, off-path government grids, and refinements of the horizon and state grid. Those checks belong to the revised calibration package. The additional checks for this companion compare the selected results with their exact records and definitions. Neither exercise generates new observations or converts a conditional numerical solution into empirical identification.

Numerical accuracy and economic scope are different questions. A small market-clearing residual indicates that a computed state satisfies the declared equations. It does not establish global uniqueness under a changed technology or validate a source-to-population transport. A finite optimizer can compare every available package allocation without proving a continuous global theorem. Refining a grid tests numerical resolution, whereas adding a different ownership or demand closure tests another model. Calling all of these exercises robustness without stating their object would obscure the evidence rather than strengthen it.

The limits are economically substantive. Contract-only extraction responses are small and negative in the selected static design, whereas combined spending stresses can reverse their sign. High incidence can apply to a very small affected flow. Unchanged output may coexist with lower profit or, in inactive configurations, with no production at either date. Covering a losing buyer can reverse the selling channel's contribution. Cobalt research support can be redundant, and low capture can eliminate government attempts despite favorable production effects. Each finding rules out a broader interpretation of the corresponding mechanism; none can be omitted merely because another scenario produces a larger contraction.

The remaining uncertainty cannot be summarized by a confidence band around simulated extrema. The chosen scenarios have no empirical probability weights, and the relevant procurement, outside-demand, innovation and intervention parameters lack a joint estimated distribution. Source measurement limitations are distinct from parameter sensitivity, model alternatives and numerical error. No assignment or sampling design presently licenses standard errors, confidence intervals or p-values for the new causal policy effects. The correct status is NOT ESTIMABLE FROM CURRENT EVIDENCE, not statistical insignificance. External-study inference remains attached to its own population and design.

The connection across models concerns incentives, not a common statistical estimand. Each specifies how an institution changes a return and which decision responds. Differences in the results show why those decisions have to be examined separately. Averaging the effects would discard the distinctions that make the comparison useful: the combined spending intervention, village charge and static contract change operate on different populations and over different horizons. Likewise, the largest simulated contraction cannot identify the preferred intervention when its instrument and implementation costs differ. Taken together, the models organize appraisal around production, incidence and continuation without supplying a single national policy score.

9 An empirical agenda tied to the next decision

The next empirical step should follow the missing links rather than seek a more precise fit to the same aggregate totals. Repeated procurement records should connect origin, assayed mineral content, transaction date, gross quote, deductions, payment timing and quantity. Candidate status and actual legal or environmental outcomes need separate measurements. The same material should be followed through certification, audits, revocations and alternative buyers. This would permit an evaluation to distinguish a cleaner register from reduced extraction and to observe the outside expenditure response on which the larger production contrasts depend.

Institutional transmission requires information about choices that ordinary sales records omit. The relevant buyer dataset includes losing offers, no-bid decisions, reserves, inventories and buyer-level coverage. Household accounts must measure essential consumption, labor, retained cash, outside income and exit options. Intermediary accounts must distinguish procurement expenses from goods-margin recovery and profit. These observations would test whether unchanged output reflects preserved livelihoods, compressed intermediary returns, a binding participation

constraint or selection among active producers. None can be inferred merely from national mineral receipts.

The dynamic and financing links call for dated cohorts. Research observations should connect spending, ownership, success, adoption and the period over which an innovation affects production. Public records should link attempted interventions to their costs, outcomes, reversals and reachable production. Tax liabilities and dated collections must use compatible bases to measure capture. Grant commitments, cash payments, fees, deposits, refunds and forfeitures need separate actor identifiers. These records would make it possible to assess a calendar financing bridge and the distribution of obligations, rather than attach an annual label to a normalized model period.

This evidence can be assembled in stages. First, the certificate must be linked to the same physical unit and transaction stage in production and buyer records. Changed market access can then be related to alternative offers and the producing household's choices. Dated adoption and public accounts are needed to interpret research and intervention over an operational horizon. The work need not take the form of one large survey. Compatible identifiers, dates and population definitions can connect records collected for different purposes. Administrative reconciliation would improve the interpretation of the existing accounts even before a design capable of identifying a causal policy effect is available.

A limited pilot with predeclared outcomes and an independently reviewed assignment or comparison design is therefore a proposed next step, not evidence of a national benefit. Its decision rules should be set before implementation, including production and livelihood measures, exposure to alternative markets, implementation costs and budget limits. The existing acquisition package contains linked empty tables and a validator; it contains no collected observations. Prospective statistical power and inference require the actual units, clusters, assignment process and observation quality. Specifying that design is a substantive research branch, not a formatting addition to the current paper.

10 Conclusion

Certification can change production when it changes the return sustaining extraction after alternative sales and institutional adjustment are taken into account. The three calibrations make that condition concrete. Better discovery can improve surviving claims while affecting extraction only slightly. Selling arrangements and household participation can absorb a restriction in profits or preserve output. Research direction and finance condition innovation-dependent adjustment, while effective capture conditions government-led conversion. Favorable market adjustment can also occur without state attempts.

The three economies show why success on one margin cannot stand in for success on another. Certificate quality does not measure production, incidence does not measure aggregate scale, and stable output does not establish welfare. Statutory rates do not measure collection, just as future fiscal returns do not supply current cash. The calibrations quantify these differences under specified assumptions and retain the cases in which the proposed channels fail. Transaction

records, household accounts and dated research and public-finance data would allow those assumptions to be confronted with observed choices. Until then, the results support conditional appraisal, without establishing causal policy effects, affordability or grounds for policy approval.

11 Global comparison of the established and nonlinear calibrations

The extension broadens the observed production and trade picture without pretending that more country labels identify unobserved supply, buyer or government behavior. Tantalum, tin and tungsten join gold and cobalt as separate calibrated mineral markets. Rough diamonds are a descriptive certification comparator; copper and nickel provide industrial and coproduct context. Physical quantities are aggregated only within a mineral.

Workstream A preserves the original empirical methods and extends their resource configurations. Workstream B reconstructs the nonlinear whole-buyer VCG and Muir economies in independent code. Both share the frozen evidence bundle. The comparison starts from the quadratic-cost benchmark, equivalent to supply exponent one; A also retains its earlier exponent and procurement alternatives. Each model and mechanism has its own anchored baseline.

12 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.

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.

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.

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.

Table 5: 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

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.

Table 6: Candidate extraction under stronger discovery: matched curvature

Mineral	Existing method, $\eta = 1$	Nonlinear Muir	Nonlinear VCG
Gold	-0.00888	-0.00666	-0.00666
Cobalt	-0.50559	-0.37916	-0.37915
Tantalum	-0.13541	-0.10155	-0.10155
Tin	-0.13504	-0.10128	-0.10128
Tungsten	-0.13207	-0.09905	-0.09905

Percentage changes from each model’s own baseline, outside spending fixed. B sends 25% of both source classes to retail; A has the original bulk construction. Unknown country shares are 0.5. These rows match curvature and observed scales, not every institutional primitive.

12.1 What is shared and what changes

The shared static result is the distinction between certified composition and physical extraction. A change in the return to claiming eligible origin can reduce surviving candidate claims much more than it changes the outside return to production. The conclusion is conditional on the source populations, contract and demand assumptions; it does not establish the effectiveness of any implemented regulation.

The nonlinear construction changes the selling institution and the processing account. Whole-owner VCG removal, all-offered inventory, Muir resale and genuine retention affect payments




Figure 1: The existing-method extension separates extraction and register effects

and actual output. Receipts divided by price cannot stand in for physical final goods. The independent bulk-only comparison isolates the retained public-price block before the retail channel is added. Country composition, route shares and technology are then separately visible sources of differences. In the displayed static reference, stronger discovery acts on the bulk contract while retail acquisition charges remain fixed. The response is therefore close to 75% of A's exponent-one response, reflecting the 25% retail route. The near-equality of Muir and VCG here does not establish general equivalence of the two institutions.

For institutional resilience, A supplies 1,120 hypothetical village allocations with zero or concave external gate retention. B supplies 216 allocations using the rebuilt selling institutions. Their fixed-route comparators are stated separately. Both calculate each policy relative to that configuration's own baseline and retain all six orders in the Shapley decomposition. Similar output resilience can coexist with different household payments, gate profit or retained inventory. Neither package has the linked household and buyer sample needed to label these as country-specific estimates. The base full-configuration VCG village exits in both paired states: its zero contraction is not preservation of positive production. Its selling switch compares nonnested VCG and posted-price menus. Some small Muir effects are close to production-grid spacing and retain the separate refinement qualification.

Dynamic differences are larger because research ownership and continuation are substantive choices. A retains the established technology and government equations across its expanded country set. B derives owner prizes from its one-period access rights, paid installation, deposits and dated refunds, excluding public adopters from private ownership returns. Its service and custody obligations have explicit funded accounts. Country governments retain their own fiscal budgets in both packages.

The B experiments are finite, with no new commitments after their terminal date and fully funded inherited and remaining liabilities. Infinite-horizon continuation remains uncertified.

Some nonconvex finite stages use a disclosed atomless allocation across genuinely indifferent optimal cells; nonattained suprema remain failures. These are limits on interpretation, not grounds to replace a failed experiment with a more favorable path. Admitted paths satisfy their finite equations and optimality conditions; neither government-path uniqueness nor selection-invariant policy effects are established.

Table 7: Period-100 candidate extraction: distinct dynamic economies

Mineral	Policy	A, $\eta = 1$	B Muir	B VCG
Gold	screening	-0.0089	-0.0067	-0.0067
Gold	combined	-13.3842	-13.3532	-13.3534
Cobalt	screening	-0.5259	-0.4008	-0.4023
Cobalt	combined	-12.5687	-12.7333	-12.7306
Tantalum	screening	–	-0.0000	-0.0000
Tantalum	combined	–	-13.2770	-13.2728
Tin	screening	-0.0002	-0.0000	0.0000
Tin	combined	-13.4110	-13.2763	-13.2720
Tungsten	screening	0.0004	0.0000	-0.0000
Tungsten	combined	-13.3157	-13.2701	–

Percentage changes from each economy’s own baseline. Both select segmented markets, supported research and active country governments, with unknown shares 0.5. B uses T120 and the separately stated small candidate retail core. Ownership, demand growth and technology differ; this is not an isolated mechanism effect. A dash is an unavailable admitted pair.


The main existing-method dynamic design admits 797 of 936 cases; a separate joint institutional-vector exercise admits 209 of 324. The remainder retain their numerical, boundary or continuation status. Only six of ten selected central exponent-.5 baselines complete all separate horizon/grid checks. This does not certify every timing result. The nonlinear dynamic ledger likewise retains each declared failure, with its own finite-horizon tests. It admits 401 of 590 cases; 189 retain an explicit numerical or admissibility reason. Its central supported comparison admits 29 of 30 paths, with tungsten–VCG combined still unresolved. Central 3T screening effects at period 100 are numerical zero after near-complete reform of the reachable population, not a general absence of earlier response. The separate T120/T140 checks cover 76 admitted pairs and cannot certify infinite continuation.

Both packages report separate country and central programme cash histories. In B, the central flow bridge excludes the initial inherited reserve and custody-provider capitalization. Private programme liquidity is measured before operating receipts and includes dated installation payments. These quantities cannot be added or relabelled as an estimated national budget.

Table 8: Classification of comparative findings

Class	Interpretation
Shared	Register composition and physical extraction require separate accounts.
Mechanism-dependent	Payment, retention, resale and ownership alter incidence and incentives.
Source-dependent	Production definitions and price stages change scales and weights.
Unresolved	Country response parameters, causal effects and general long-run continuation.

Agreement between related conditional models is not independent empirical confirmation.



2024 mine content times the admitted reference price; product stages differ.
Weights describe an accounting index, not actual producer revenue or welfare.

Figure 2: The fixed-reference-value index and its concentration

Table 9: Separate 2025 update: USGS world mine-production estimates

Mineral	2024, same USGS vintage	2025 estimate	Change (%)
Gold	3,280	3,300	0.61
Cobalt	302,000	310,000	2.65
Tantalum	2,500	2,500	0.00
Tin	294,000	290,000	-1.36
Tungsten	82,000	85,000	3.66

Metric tonnes of metal content; rounded world totals from Mineral Commodity Summaries, February 2026, commodity-table page 2. These supplementary estimates do not update the calibration. USGS gold is a separate source scenario from the preserved WGC 2024 scale; no cross-source growth rate is calculated.

12.2 Evidence priorities and publication claims

A matched procurement panel needs assays, payable content, deductions, credit, food and other in-kind payments, transport and the date-specific benchmark. Those observations must be linked to a defined mining population and actual production costs. The Rwanda tantalum table demonstrates why a quoted price ratio cannot be admitted before its units and arithmetic are reconciled.


Effective exposure needs dated covered buyers and products, re-export routes, actual enforcement and feasible outside sales. Innovation requires ownership and installation terms, researcher payments and subsequent output, not an aggregate R&D ratio alone. State-capacity calibration requires actual commodity-specific collections, protection and intervention costs. The EITI and mine-site records supply context and sampling leads, not those parameters.

The published numerical claims trace to saved equilibrium objects and their source or assumption records. A failure remains in the experiment denominator. Conditional ranges are not confidence limits, and no p-value is manufactured from a scenario grid. This allows a stronger empirical design to replace an assumption later without obscuring which earlier conclusion depended on it.

12.3 Primary source access

The expanded source archive records individual pages, request parameters and hashes for USGS commodity statistics, BGS world mineral production, UN Comtrade, the Rwanda 3T study, IPIS visits, EITI workbooks, World Bank prices and Kimberley Process statistics. Primary portals are <https://www.usgs.gov/centers/national-minerals-information-center/commodity-statistics-and-information>, <https://www.bgs.ac.uk/mineralsuk/statistics/world-mineral-statistics/>, <https://comtradeplus.un.org/> and https://kimberleyprocessstatistics.org/public_statistics. The companion papers give the detailed source-specific citations and audits.

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 3: Global descriptive production, separate BGS source: first four resources

Worldwide BGS production records, 2024 | Part 2




Export proxies (hatched): Tungsten: Rwanda. Diamonds: Ghana.

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 descriptive production, separate BGS source: remaining resources

Observed customs destination concentration




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

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

Figure 5: Reported destination concentration

Selected six-origin 3T routes to China



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

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

Figure 6: Selected-origin routing and coverage

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 7: Evidence coverage and limits of calibration

Data and code availability

The versioned packages are `global_calibration_extension` and `consolidated_calibration`. They contain the frozen source and assumption registers, original benchmark snapshot, shared 2024 registry, new public-source panels, full experiment registers, reproducible code and independent checks. All original files are preserved. The common schema identifies model, mechanism, mineral, country, source scenario, policy, comparator, horizon and outcome. Publications are generated from those saved results; unavailable pairs are not imputed. Public source locators, archive hashes and acquisition limitations are recorded. No public repository identifier or journal acceptance is claimed.

A Static certification and procurement

The following definitions reproduce the maintained revised calibration, with notation local to this appendix. They specify the numerical claims in Table 1; they do not extend the parent’s common-curvature theorem to every alternative technology. 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

$$A = r\ell_L - \ell_E, \quad a = A/\kappa, \quad (2)$$

$$\pi = \frac{1 - x - rd}{c + d(1 - r)a}, \quad e = a\pi, \quad p = r + (1 - r)e. \quad (3)$$

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 $p_{D0} = x_0$.

The complete static grid crosses common curvature $\eta_0 \in \{.25, .50, 1\}$; initial quote $x_0 \in \{.70, .80, .85\}$; outside-spending multiplier $\lambda_D \in \{.80, 1, 1.20\}$; and policy $(r, \kappa) \in \{.35, .50\} \times \{.25, .50\}$. Gold includes common-curvature and heterogeneous branches, each with mining alone or fixed recycling. In the heterogeneous branch, Colombia uses $\zeta_C = .70$, Peru uses $\zeta_P \in \{.85, .875, .90\}$, and net recovery $m \in \{0, .10, .20\}$ enters equation (1); other supply retains η_0 . The same recovery-grid value is used in the declared joint regional vector. Cobalt has single-metal and fixed-mix joint-revenue branches. Baseline late discovery and diligence are $(.25, .50)$. Contract-only comparisons use reference and changed contracts at the same λ_D ; combined comparisons use the changed contract and spending against reference contract and $\lambda_D = 1$. These are finite-design choices, not probability draws.

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

$$C_i(h) = C_{i0}[1 - w + wh]^{\eta_{C_i}}, \quad D_i(p_D) = D_{i0}[1 - w + w(p_D/x_0)]^{\eta_{D_i}}. \quad (4)$$

Mining-only gold and single-metal cobalt set $w = 1$; joint-revenue cobalt sets $w \simeq 0.1728$. 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

$$\mathcal{R}_i = r\ell_L\pi_i - \frac{A^2\pi_i^2}{2\kappa}. \quad (5)$$

The two market equations are

$$E_C = h \sum_i (1 + \phi + \mathcal{R}_i) X_i, \quad E_D = xh \sum_i Y_i. \quad (6)$$

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

$$X = \frac{C}{1 - \pi}, \quad Y = D - \pi X > 0, \quad (7)$$

$$\text{revoked} = p\pi X, \quad \text{live candidate} = (1 - p)\pi X. \quad (8)$$

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

$$\frac{K(q)}{Pq} = \frac{1}{1 + 1/\eta} = \frac{\eta}{1 + \eta}. \quad (9)$$

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

$$\eta = \frac{\zeta - m}{1 - \zeta + m}, \quad 0 < \zeta - m < 1. \quad (10)$$

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,

$$E_D = xhY \leq xhD(xh) < p_{\max}D(p_{\max}). \quad (11)$$

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 Transmission and the household economy

B.1 The household and intermediary

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

$$\max_{\ell, x, b} U = \log(t + x - c) + \eta_h \log(1 - \ell) + \beta_h \log(b + b_0), \quad (12)$$

$$(b_z + \mu)x + b = a + \theta A\ell, \quad t + x > c. \quad (13)$$

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

$$U \geq \bar{U} = \log s_o + \eta_h \log L_o + \beta_h \log(b_o + b_0), \quad (14)$$

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

Table 10: Assumed base primitives of the village experiment

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

The three switches are explicit. With interface I absent, $\mu = 0$; with I present, μ is chosen from 41 equally spaced quotes on $[0, \bar{\mu}]$. With subsistence dependence H present, $t = 0$; without

H , $t = c$ is delivered independently while a is unchanged. Selling S selects the Muir or VCG mechanism specified below instead of that buyer economy's fixed contracts. Removing H therefore changes an outside resource, not merely a market label.

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

$$w_{b=0} = \frac{\eta_h[a - B(c - t)]}{1 - (1 + \eta_h)\ell}, \quad (15)$$

$$w_{b>0} = \frac{\eta_h[a - B(c - t) + b_0]}{(1 + \beta_h) - (1 + \beta_h + \eta_h)\ell}. \quad (16)$$

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

The gate's attained implementation cost is

$$C_{IH}(q) = \min_{\text{feasible supporting terms}} \{\theta q - \mu x + \mathcal{T}(q)\}. \quad (17)$$

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

B.2 Muir receipts and the fixed-route comparator

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

$$P_U(u) = \begin{cases} 1 - 8u/3, & 0 \leq u \leq 1/4, \\ 4(1 - u)/9, & 1/4 < u \leq 1. \end{cases} \quad (18)$$

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

$$v(u) = \begin{cases} h_M + s_M u, & a_M \leq u \leq b_M, \\ u P_U(u), & \text{otherwise.} \end{cases} \quad (19)$$

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

$$W_M(q, \tau) = \begin{cases} (M - \tau)q, & 0 \leq q \leq R, \\ (1 - \rho_M)\{v(q - R) + R(M - \tau)\} + \rho_M q P_U(q - R), & R < q \leq 1 + R. \end{cases} \quad (20)$$

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

$$F_M(q, \tau) = \kappa_M q(M - \tau) + (1 - \kappa_M)qP_U((1 - \kappa_M)q), \quad (21)$$

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

B.3 Finite whole-buyer VCG and fixed contracts

VCG uses a different buyer economy: buyers $i = 0, 1, 2$ have gross marginal values $g_i(z) = (1.8 - 2z, 0.9 - z, 0.75 - 0.75z)_i$ for $0 \leq z \leq 0.4$. Physical packages have mass $\Delta = 2.5 \times 10^{-5}$. Package j is valued at its midpoint $z_j = (j - 1/2)\Delta$, so its gross bid is $b_{ij} = \Delta g_i(z_j)$. Let $f_i \in \{0, 1\}$ indicate whole-buyer coverage, with $(f_0, f_1, f_2) = (1, 0, 0)$ in the base case. Net cumulative value for k packages is

$$V_i(k; \tau) = \sum_{j=1}^k (b_{ij} - \Delta \tau f_i). \quad (22)$$

For an offered capacity of n packages, calculate

$$\mathcal{W}(n) = \max_{\substack{k_i \in \mathbb{Z}_{\geq 0} \\ \sum_i k_i \leq n}} \sum_i V_i(k_i; \tau), \quad \mathcal{W}_{-i}(n) = \max_{\substack{k_i=0, k_h \in \mathbb{Z}_{\geq 0} \\ \sum_h k_h \leq n}} \sum_{h \neq i} V_h(k_h; \tau). \quad (23)$$

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

$$T_i(n) = \mathcal{W}_{-i}(n) - [\mathcal{W}(n) - V_i(k_i^*(n); \tau)]. \quad (24)$$

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

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

B.4 Declared sensitivity and finite scope

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

Table 11: Added joint primitive bundles

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

The reproduced source set additionally changes one primitive at a time: halves the charge, sets wealth to 0.01 or 0.03, halves the margin cap, doubles the Muir fixed share and the VCG buyer-0 fixed share, or covers buyer 1 in VCG. Its refinement doubles Muir quantity resolution and goods-margin resolution and halves VCG package mass. Every optimum is certified only over its stated finite quantity, margin and capacity alternatives, with the exact household branches above. There is no empirical weighting of these bundles and no claim that finite-grid endpoints bound all intermediate primitives.

B.5 Incidence and channel accounting

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

$$\frac{p_1}{p_0} = (1 - \delta)^{1/\eta}, \quad \frac{O_1}{O_0} = (1 - \delta)^{-1/\eta}, \quad (25)$$

$$E_0 - E_1 = (Q_0 - Q_1) + (O_1 - O_0) \quad (26)$$

$$= Q_0\delta + O_0[(1 - \delta)^{-1/\eta} - 1]. \quad (27)$$

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

$$T = \frac{Q_0 - Q_1}{E_0 - E_1} = \frac{\delta}{\delta + (O_0/Q_0)[(1 - \delta)^{-1/\eta} - 1]}. \quad (28)$$

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

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

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

$$\varphi_i = \sum_{J \subseteq N \setminus \{i\}} \frac{|J|!(2 - |J|)!}{3!} [R(J) - R(J \cup \{i\})]. \quad (29)$$

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

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

$$\theta q - \mu x + \mathcal{T}(q) = b_z x + b - a + \mathcal{T}(q). \quad (30)$$

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

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

$$W(q, \tau) = (1 - \rho)[v(q - R) + R(M - \tau)] + \rho q P_U(q - R). \quad (31)$$

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

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

$$\Delta \sum_i w_i s_i = \sum_i \frac{w_{i1} + w_{i0}}{2} \Delta s_i + \sum_i \frac{s_{i1} + s_{i0}}{2} \Delta w_i. \quad (32)$$

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

C Dynamic adjustment and finance

C.1 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 12: 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 $\chi = .50$

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.

C.2 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

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

Initial technologies satisfy $A_{C0} = A_{D0} = 1$, and each focal country has $m_{j0} = 1$. The initial public vintage is $A_{C,-1} = 1/g_{ref}$. Quantities are normalized by world mine output. The candidate coefficient is $B_{Dj} = D_{j0}/x_0^\eta$; eligible coefficients split initial eligible production between safe and vulnerable capacity using the declared .10 focal vulnerable share. The remaining world eligible share is .50. These normalizations fix the initial fiscal comparison as well as production.

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.3 Research, spending fit and technology timing

Research owners take the realized equilibrium prices and prizes as given when allocating effort; they do not internalize derivatives of aggregate prices with respect to their own research. 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

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

Define the owned eligible invoice as $I_{Ct}^{\text{owned}} = B_s A_{Ct} h_t^{1+\eta}$, its profit as $\Pi_{Ct}^{\text{owned}} = I_{Ct}^{\text{owned}} / (1 + \eta)$, and total candidate profit as $\Pi_{Dt} = p_{Dt} Q_{Dt} / (1 + \eta)$. These definitions exclude publicly converted eligible capacity from the privately owned research prize.

Let $v_C = 1 + \gamma z(s_C)$, $v_D = 1 + \gamma z(s_D)$, and laboratory curvature $\ell = .002$, 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. The increment breaks numerical ties and is negligible at the reported spending scale; it does not represent an estimated 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

$$w \cdot 1 + \frac{.002}{2}(s_C^2 + s_D^2) + \frac{.002}{2}(.20)^2 = r_{RD}(h_0 Q_{C0} + p_{D0} Q_{D0}), \quad s_C + s_D = .80.$$

The research allocation and prizes in this fit use the reference fixed-expenditure next-period experiment, not the later growing-demand period-zero path. The central conditional fit gives $\gamma \simeq 3.773 \times 10^{-2}$ and $g \simeq 1.01606$ for gold; the saved metadata retain full precision for every commodity/vector. Equivalently, 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, fitted technology parameter and assumed growth process are therefore distinct objects.

C.4 Dynamic government comparison and local frontiers

The four policy vectors are central $(\chi, \beta, \rho, r_{RD}, x_0) = (1, .25, .25, .005, .80)$; low capture $(.50, .25, .25, .005, .80)$; constrained $(.50, .15, .10, .0025, .85)$; and enabling $(1, .40, .30, .007417, .70)$. Each is evaluated under private and supported research and baseline, screening and combined policy, with other primitives fixed within each pair. Screening sets $r = .50$; baseline $r = .25$; combined also sets outside-expenditure response $q = .25$. None of these joint-vector comparisons identifies a separate causal effect of each changed component.

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

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

The numerical backward recursion sets $U_T = V_T = 0$ at its terminal truncation date. The reported horizons and separate refinements below evaluate the sensitivity of the reporting interval to that finite approximation. The computed paths are not an infinite-horizon existence or uniqueness theorem.

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 10^{-5} , the archived reporting convention. The underlying complementarity check retains actual rates and its stated boundary tolerances.

C.5 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

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

The complete normalized operating bill C_t is service plus verification research plus verification deployment. Let n_t be new certified admissions (X_t under segmentation and zero under pooling), o_t the sum of the preceding H admissions cohorts, K_t the accumulated record count, and v_t the vulnerable-origin invoice ($h_t B_v h_t^\eta / (1 - \pi_t)$ under segmentation, zero under pooling). With $Z_t = (g_{ref} 1.005)^t$, the assumed service bill is

$$S_t = .0015 + .005n_t/Z_t + .002o_t/Z_t + .0002K_t/Z_t + .01v_t.$$

Verification research costs $w_t(.20) + (.002/2)(.20)^2 L_t$, and deployment costs $.0002L_t z(.20)$. Therefore C_t is their sum with S_t . Initially $K_0 = 5X_0$ and the previous H admissions equal X_0 ; each period appends its new admissions to the cohort history and accumulated records. The trust begins with H baseline deposit cohorts, each equal to d times the baseline vulnerable invoice plus b_o times the baseline safe invoice, with corresponding baseline expected forfeitures. Matured cohorts settle before new cohorts are added. These are assumed operating technologies and expected accounts, not observed administrative unit costs.

Dollar accounts multiply normalized values by $V_0 = Q_{world,0} P_0$, with $Q_{world,0}$ in tonnes. Gold uses $P_0 = 2386.2 \times 10^6 / 31.1034768$ USD per tonne (the troy-ounce conversion); cobalt uses $P_0 = 11.84 \times 1000 / .45359237$ USD per tonne (the pound conversion). Discounted totals sum the defined flow over periods 0–100 with factor $.95^t$. This conversion supplies nominal reference-year scale, not calendar inflation or an annual budget.

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 $A_t(c) = \max\{0, C_t + G_t^{\text{paid}} + cB_t - F_t\}$, 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.

C.6 Paired effects, explored ranges and inference

For each parameter 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 too small to support a meaningful percentage, the outcome is a signed government net budget balance, or a percentage-point difference is the defined unit. Signed fiscal balances are reported as dollar differences. 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.

C.7 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 cases with government intervention enabled 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.

The principal runs use $T = 500$, $N = 301$ and fixed-point tolerance 2×10^{-8} . The enabling vector uses $T = 700$, $N = 501$. Frozen segmented runs without government intervention use $T = 110$ because they have no forward government value problem and no terminal government continuation is required. 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 10^{-8} , and dynamic complementarity and on-path Bellman checks use 10^{-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. Conservative rounded upper bounds on the initial discrepancies are 2.78×10^{-13} for outside clearing, 2.47×10^{-14} for research KKT conditions, 1.26×10^{-9} for government complementarity and 1.68×10^{-9} for the Bellman difference. 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.

D Numerical exhibit crosswalk and reproduction scope

All paths below are relative to `calibration_revision_2026-09-24/`. The three revised manuscripts remain separately available [Lucio, 2026a]. This companion reorganizes their verified results without generating additional empirical observations. The publication claim

register preserves full-precision values, row/scenario selectors, denominators and conditions.

Table 13: Displayed results and exact records

Exhibit	Exact result and selectors	Existing verification
Table 1	<code>results/c1/headline_candidate_ranges.csv</code> ; Peru/Colombia heterogeneous mining; DRC single/joint revenue; contract and combined	<code>checks/C1_VERIFICATION_CERTIFICATE.md</code>
Table 2	<code>results/c2/screening_policy_effects.csv</code> ; gold/cobalt reference; gold Peru upper procurement transport	<code>checks/C2_INDEPENDENT_CHECKS.json</code>
Table 3	<code>results/c3/effect_ranges.csv</code> ; period 100 world candidate output; private/supported, screening/combined against baseline	<code>checks/C3_INDEPENDENT_C1_CHECKS.json</code>
Table 4	<code>results/c3/policy_financing_frontier.csv</code> ; central gold supported, combined versus baseline; grant and fee-net deficit PV	Same C3 finance reconstruction; equation checks separately retained

Village levels use `results/c2/institutional_paired_effects.csv`, scenario base; institutions retain their own no-charge comparator. Channel findings use `results/c2/institutional_shapley.csv`. The trade decomposition is `results/c2/cobalt_composition_decomposition.json`, retaining the separate population bridge. First-attempt and capture comparisons use C3 scenario/path records and paired summaries. Source scales and assumptions remain in the admission and identification ledgers.

The new publication checks reconstruct selected saved-result arithmetic and compare editions with a common claim register. They are separate from earlier solver, corruption, accounting and refinement checks. Assertion counts are not independent observations. Full reproduction requires the pinned numerical runtime documented in `REPRODUCTION.md`. Building this paper requires LaTeX and the supplied sources and bibliography; it does not require regenerating calibrations.

Appendices A–C give the equations, primitives, domains, policy/comparator definitions and accounts needed to assess the central claims. Exhaustive scenarios, logs and off-path grids remain in the supplied research materials. Source-access restrictions, especially provisional Peruvian procurement, remain binding even where the arithmetic is exact. The version register identifies the parent source tree and its matching PDF.

References

- Daron Acemoglu. Directed technical change. *Review of Economic Studies*, 69(4):781–809, 2002. URL https://www.nber.org/system/files/working_papers/w8287/w8287.pdf. Accessible version: NBER Working Paper 8287, May 2001.
- Daron Acemoglu, Simon Johnson, and James A. Robinson. Institutions as a fundamental cause of long-run growth. In Philippe Aghion and Steven N. Durlauf, editors, *Handbook of Economic Growth*, volume 1A, pages 385–472. Elsevier, 2005. doi: 10.1016/S1574-0684(05)01006-3. URL <https://economics.mit.edu/sites/default/files/publications/institutions-as-the-fundamental-cause-of-long-run-.pdf>.
- ANM–UPME. Boletín estadístico informativo minería en cifras. Technical Report 24, ANM–UPME, 2025. URL https://docs.upme.gov.co/SIMEC/SIMCO/Cifras-Sectoriales/BoletinMinerales/Boletin_24_mineria_en_cifras_diciembre_2025.pdf. December, page 7. Frozen preliminary 2024 royalty-associated production vintage.
- Banco Central de Reserva del Perú. Memoria 2025. Technical report, BCRP, 2026. URL <https://www.bcrp.gob.pe/docs/Publicaciones/Memoria/2025/memoria-bcrp-2025-1.pdf>. Indexed commercialization passage. Provisional: archived nominal PDF is HTML; underlying USAID observation date and sample unresolved.
- Timothy Besley and Torsten Persson. The origins of state capacity: Property rights, taxation, and politics. *American Economic Review*, 99(4):1218–1244, 2009. doi: 10.1257/aer.99.4.1218. URL <https://www.aeaweb.org/articles?id=10.1257/aer.99.4.1218>.
- CMOC Group Limited. Annual report 2024. Technical report, CMOC, 2025. URL <https://www1.hkexnews.hk/listedco/listconews/sehk/2025/0424/2025042400242.pdf>. Page 6 (PDF page 7). DRC copper/cobalt segment revenue; industrial comparator.
- Democratic Republic of Congo. Revised mining code, 2018. URL <https://www.leganet.cd/Legislation/J0/2018/JOS.28.03.2018.pdf>. Article 247, 30% profits-tax anchor. Official legal text mirrored by Leganet.
- DIAN. Renta personas jurídicas, 2026. URL <https://www.dian.gov.co/impuestos/sociedades/Paginas/Renta-personas-juridicas.aspx>. Frozen source accessed 24 September 2026. General 35% anchor; not the complete mining fiscal regime.
- Andrew L. Gulley. 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, 2023. doi: 10.1073/pnas.2212037120. URL <https://doi.org/10.1073/pnas.2212037120>.
- Simon Loertscher and Ellen V. Muir. Monopoly pricing, optimal randomization, and resale. *Journal of Political Economy*, 130(3):566–635, 2022. doi: 10.1086/717894. URL <https://www.journals.uchicago.edu/doi/10.1086/717894>. Publisher metadata and abstract verified; author working paper of 19 December 2019 also consulted.
- Elies Lucio. Revised calibration package: certification, transmission and state capacity. 24 September. Three preserved calibration manuscripts, exact results and verification records. Local materials; no public repository identifier assigned, 2026a.
- Elies Lucio. For a Handful of Gold. Manuscript, editorial revision of 13 September 2026. Verified 134-page root PDF and included source tree, 2026b.

- Ministerio de Energía y Minas del Perú. Anuario minero 2024. Technical report, MINEM, 2025. URL <https://cdn.www.gob.pe/uploads/document/file/8155729/6827926-anuario-minero-2024-minem.pdf?v=1749053631>. Pages 66–67. Fine-gold production and administrative strata.
- OECD. Enhancing regional mining ecosystems in the european union. Technical report, OECD Publishing, 2025. URL <https://doi.org/10.1787/97ba1224-en>. Box 3.9, page 96. Broad mining research-spending benchmark.
- Dominic P. Parker and Bryan Vadheim. Resource cursed or policy cursed? US regulation of conflict minerals and violence in the Congo. *Journal of the Association of Environmental and Resource Economists*, 4(1):1–49, 2017. doi: 10.1086/689865. URL <https://www.journals.uchicago.edu/doi/10.1086/689865>.
- Rio Tinto. Annual report 2024. Technical report, Rio Tinto, 2025. URL <https://www.riotinto.com/-/media/content/documents/invest/reports/annual-reports/2024-annual-report.pdf?rev=af45e9c438764f07ab7944c263ca3615>. Page 178, notes 6–7. Group research and consolidated sales; transported spending comparator.
- SUNAT. Tasa del impuesto: Régimen general, 2026. URL <https://orientacion.sunat.gob.pe/03-tasa-del-impuesto>. Frozen source accessed 24 September 2026. General 29.5% anchor; not effective collection.
- United Nations Statistics Division. UN Comtrade database: archived mineral-trade extraction, 2026. URL <https://comtradeplus.un.org/>. Archived extraction, not refreshed. Exact response URLs and hashes in work/c2/archived_inputs/manifest.json of the supplied calibration package.
- UNODC and Ministerio de Minas y Energía. Colombia: Explotación de oro de aluvión. evidencias a partir de percepción remota 2023. Technical report, UNODC and MME, 2025. URL https://biesimci.org/fileadmin/2025/documentos/evidencias_de_explotacion_de_oro_de_aluvion_2023.pdf. Procurement page 50; production categories Table 5, page 56.
- U.S. Geological Survey. Mineral commodity summaries: Cobalt. Technical report, U.S. Geological Survey, 2026. URL <https://pubs.usgs.gov/periodicals/mcs2026/mcs2026-cobalt.pdf>. Pages 1–2. Revised 2024 production and refined-price vector.
- World Gold Council. Gold demand trends: Full year 2024. Technical report, World Gold Council, 2025. URL <https://www.gold.org/goldhub/research/gold-demand-trends/gold-demand-trends-full-year-2024>. 5 February. Supply and annual reference-price tables.