

For a Fistful of Gold

Elies Lucio *

4 October 2026

Abstract

I develop a theory of how the organization of local markets shapes the effects of supply-chain regulation on illegal mining. An armed intermediary buys gold from households and sells them essential imports. What it pays for gold, it recovers through import margins. Households must cover their subsistence needs. Falling returns can make them mine more, not less. Downstream, laundering, unregulated buyers, and retention give the intermediary ways around restrictions. Restrictions can therefore change sales and receipts without changing extraction. Imperfect certification instead separates certified from unverified gold. This raises the profitability and fiscal value of legal mining. Once funded, clean innovation expands legal opportunities and the prospective tax base. Authorities who profit from illegal mining then find intervention optimal once the fiscal value of legal mining is high enough. This holds even though intervention is costly, may fail, and can be postponed. With livelihoods financed and production safe, illegal mining can end.

Keywords: illegal mining; intermediation; mechanism design; provenance; directed technical change; state capacity; conservation.

*Université Paris-Dauphine–PSL. The author thanks Gilles Chemla, Joseph Kaboski, Philippe Aghion, Simon Johnson, Lars Peter Hansen, Fernando Alvarez, Catherine Casamatta, Quentin Vanderweyer, Jerome Dugast, Olivier Tercieux, Evgenia Passari, Olivier Darmouni, Stephen Duffie for their comments and guidance. He also thanks the participants in internal seminars at Notre Dame, Paris-Dauphine and Collège de France, and participants at the LAGV symposium, ISCFS, the Paris-Dauphine CEPR PhD workshop, MFR Chicago and the HEC Paris PhD workshop for their insights and companionship.

1 Introduction

A kilogram of gold refined from a licensed mine and a kilogram refined from a pit held by an armed group are the same object. Origin is not an attribute of the metal but a claim about a transaction, and in a market that trades some \$361 billion of gold a day ([World Gold Council, 2026](#)), a claim is easy to place. Peru shows what follows. Armed groups now fight licensed producers for veins inside their own concessions ([Bloomberg News, 2025](#)), and in the first half of 2026 illegally sourced gold made up 54 percent of the volume of the country's gold exports ([Samon Ros, 2026](#)). That gold did not reach the world market for lack of rules. Since 2010, U.S. conflict-minerals disclosure, the OECD due diligence guidance and its gold supplement, the London Bullion Market Association's responsible-sourcing program, and a binding European regulation in force since 2021 have each asked the buyers of gold to vouch for its origin ([SEC, 2012](#); [OECD, 2016](#); [LBMA, 2012](#); [European Union, 2017](#)). Their costs fall on the firms that comply: the U.S. rule alone was expected to cost issuers \$3 to \$4 billion at the outset ([SEC, 2012](#)). Their reach stops where the claim is written.

Behind the claim lie costs that no buyer pays. Across Africa, the commodity boom of 1997–2010 may account for as much as a quarter of the average level of violence ([Berman et al., 2017](#)). In the Amazon, the damage is on the intensive margin: mining does not push a frontier forward but opens the forest from within, patch by patch, in a biome that may already be near a critical fragmentation threshold ([Saravia et al., 2018](#)). Degraded forest can lose up to a third of its dry-season evapotranspiration, and with it part of the moisture the winds carry toward the Andes ([Lapola et al., 2023](#); [Ruiz-Vásquez et al., 2020](#)). What these places share is distance from the state: the road between them and the world runs through an armed intermediary. At Rubaya, which supplies over 15 percent of the world's tantalum, that intermediary is the M23 ([Banchereau, 2024](#); [United Nations Group of Experts, 2024](#)). This paper is about that gatekeeper, and about why the price he commands survives the rules written to undercut it.

Global policies against illegal gold assume the Walrasian chain: demand moves the price, and the price moves supply. Illegal gold has no such price. Its market is fragmented, and every price between the buyer and the pit is set by someone with a stake in it. In Section 2 we study, in a stylized region we call Selva, why extraction is resilient in that market, in the sense that it persists under regulation. The gatekeeper is the region's monopsonist for gold and its monopolist for everything gold buys. A fall in his receipts therefore need not pass to the miners at all. Holding both prices, the choice of villages and the means of coercion, he plays social planner for the region, a stationary bandit in the sense of [Sánchez de la Sierra \(2020\)](#), with his own profit as the welfare function. Because many combinations of prices deliver the same output at the same cost, a fall in what the gold is worth to him can come out of his margin rather than out of the ground. And where families cannot cover their necessities, a lower price for gold buys more digging, not less. [Braverman and Stiglitz \(1982\)](#) showed that terms in one market shape returns in another, and [Kranton and Minehart \(2001\)](#) and [Manea \(2018, 2021\)](#) that the structure of a trading

network decides who captures its surplus; we combine the two at an armed chokepoint and solve for the allocation the gatekeeper can implement inside the region. Yet none of this may be needed. Before any price reaches the miners, the gatekeeper must sell the gold, and the way he sells it can move the metal to buyers the rules do not reach, who then hold up his price. Section 3 studies that sale.

The sale is where the price survives. The gatekeeper sells through a Vickrey–Clarke–Groves auction in which he may keep any unit. His price rests on two buffers. One is his own reservation value: he can always keep the gold and launder it, so his laundering profit is a bid in his own auction. Laundering lets the gold escape the regulation unit by unit, until the margin on laundering one more unit falls to zero, if it ever does. The other buffer is competition: buyers outside the rules keep bidding, and their demand holds the price once laundering is exhausted. A tax on compliant buyers moves neither. It reshuffles who buys and what they pay; where the buffers price the marginal kilogram, output does not move. Krishna and Perry (1998) show that this auction maximizes revenue among efficient mechanisms, and Ausubel and Cramton (1999) that efficiency is optimal when resale is perfect; we extend both to a seller who can keep the good, and show that the auction is the best the gatekeeper can do within a broad class under conditions we state.

The evidence fits: rules on buyers moved the conflict rather than ending it. After U.S. conflict-mineral rules took effect in eastern Congo, militia battles shifted toward gold territories (Parker and Vadheim, 2017). A policy that bites must change what is being sold, not who may buy it. Section 4 studies the institution that would do so: a blockchain ledger that binds each origin claim to physically identified gold and preserves its history after the sale, so that certification can be withdrawn without erasing the record. The ledger alone is not enough. As Halaburda (2018) emphasizes, the functions attributed to blockchain must be distinguished from the ledger itself; here, credibility rests on evidence, physical custody and the incentive to act on them. We model a claim whose liability outlives the first sale, building on costly verification (Townsend, 1979), collusion with verifiers (Tirole, 1986) and evidence that auditor incentives shape reporting (Duflo et al., 2013). A fraud then pays only if its gain exceeds the funded cost of buying admission and, where needed, silence. Some fraud remains profitable, and some illegal gold leaks into certified stock. Withdrawing discovered claims lowers that leakage below its rate at admission, though not to zero unless every fraud is caught, and the scrutiny it requires is a recurring cost that seizures alone cannot finance. A cleaner certified stock is also compatible with continued extraction: gold kept out of certified trade can still be sold elsewhere, and Sections 5–7 follow it into the world market.

In the world market, verification splits gold in two, and the gap between the certified price and the discount on unverified gold works like a scissor with two blades (Section 6). The first blade is global. Investors go where profits are (Acemoglu, 2002; Acemoglu et al., 2012), and once verification rewards clean supply, research and certification both find funding (Section 8). The second blade is local, and it is where the state comes back. Selva also has a local authority, and the illegal trade pays it to look away: protection is income. But as clean supply expands, so does the tax base the authority would gain


by taking control ([Besley and Persson, 2009](#)), while what the illegal trade can pay stays bounded by the sites that remain, even when reform elsewhere makes illegal gold scarcer and lifts their rents for a while. Sooner or later the arithmetic turns. After a finite date, every authority that has not yet taken control finds it worth paying to do so, even though it may fail and could wait. The illegal frontier then shrinks period after period, and, with legal mining kept on safe sites, disturbance in the protected forest fades toward zero. The authority whose income depended on suppressing the state comes to maximize it by building one.

The policy implications follow directly. Regulating who may buy gold moves the metal, not the pit. What moves the pit is a record bound to the unit at extraction, scrutiny funded well enough to keep that record true, and a legal opportunity that grows faster than what protection can pay. Enforcement itself is not the obstacle: where the state reaches, satellite monitoring combined with enforcement has reduced deforestation in Brazil ([Assunção et al., 2023](#)). The obstacle is reach, and the paper shows how markets can make extending it worth a government's while. Timing still matters. Because forest loss is irreversible, conservation requires control to arrive within the damage the forest can absorb, not merely eventually (Section 8). And where protection remains as attractive as the legal flow, or intervention cannot succeed, waiting forever stays optimal; these conditions are precisely what policy must change.

Gold is only the sharpest case. Wherever a market reaches past the edge of the state, the old chain from demand to price to supply breaks, and what replaces it is a set of institutions (a gatekeeper's auction, a ledger, a government's choice) that can swallow a policy, send it elsewhere or turn it around. In those markets, what is sold is a claim, and the claim is what policy must change. And the same institutions that make a claim true can make the state worth building, so that trade, which so often finances the absence of the state, can pay for its arrival. The same questions hang over coltan, timber, cattle and fish, everywhere origin is a story told about a good rather than a fact written into it. A kilogram of gold carries no memory of where it was dug. The institutions that can give it one turn out to be the institutions that build states.

Mining and governance: selected observations in 2023

Different measures; percentages are not national illegal-output rates.



2 Selva: production behind an armed chokepoint




Figure 2: Commodity access and institutional reach: exterior mineral and import routes cross the armed gate, while public enforcement initially reaches exterior actors; gold frames mark the section’s focus and muted elements retain context.

2.1 Market access and the reach of public authority

Households in Selva depend on the mineral trade for goods they cannot produce. They grow food and exchange local products, but mining supplies the earnings with which they buy imported necessities and productive inputs. An armed intermediary controls their access to outside markets. It purchases exports and supplies imports, choosing both the return to mining and the terms on which mining income can be spent.

Joint control of these transactions is the source of the intermediary’s economic power. The theory of interlinked markets emphasizes that terms in one market affect the returns available in another [Braverman and Stiglitz, 1982]. Selva places this interdependence within a trading network whose external access is controlled by force. We derive the household allocations that the intermediary can induce and the cost of delivering mineral quantity q . These choices determine whether a change in demand outside the region reaches production inside it.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{L})$ be the directed network of usable trade routes, with villages and production sites in $\mathcal{J} \subset \mathcal{V}$ and the world market outside it. Usability depends on institutions as well as geography. A seized shipment, an impassable crossing, or unavailable credit can close a route. This network describes trade under armed control. Public intervention can subsequently replace the regime that governs access.

Definition 2.1 (A regional chokepoint). The region has a chokepoint g if every admissible path carrying mineral from a village in $\mathcal{J}$ to an outside buyer passes through g . The gate is bilateral if every admissible inward path for necessities and mining inputs also

passes through g . The armed group controls access to g and can enforce its posted terms there. Villages have no accessible alternative mineral buyer or import supplier on identical delivery and financing terms.

Removing the gate disconnects the relevant external trade paths: it is a directed articulation point. Farmers, miners and traders can nevertheless exchange goods within the region. External monopoly and internal specialization therefore coexist. As in the economics of intermediation networks, the available paths determine which traders can command rents [Manea, 2018]; here, armed control sustains the bottleneck through which the region trades.

Lemma 2.2 (Trade must cross the gate). *Suppose flow is conserved at intermediate nodes and the gate is the only external route. Every positive outward delivery of regional mineral uses gate capacity. If the gate controls both directions, every positive inward delivery of the specified imported goods uses that capacity as well.*

Proof. Decompose a feasible nonnegative flow into source–sink paths and cycles. Cycles do not deliver goods from a regional source to an outside sink. Each source–sink path must contain g by Definition 2.1, so its delivered flow is bounded by the capacity made available at that node. Apply the same argument to the directed inward network for imports. $\square$

Trade and public authority have different reach. Formal buyers, financial intermediaries and importers can face enforceable obligations while the armed organization governing the interior remains beyond direct public control. This distinction determines the initial set of policy instruments.

Assumption 2.3 (Enforcement reach). During the initial policy experiment, the global policymaker can impose obligations on a specified set of downstream agents. It cannot directly enforce extraction limits, labor rules, procurement contracts or transfers inside the armed-controlled region. Local production technology, the gate’s enforceable control and the regional outside opportunities are fixed in that experiment. Establishing public authority inside the region requires a distinct, costly institutional intervention.

The intermediary consequently decides how global demand affects village incentives. An exterior policy changes its opportunities, and the intermediary responds by choosing the local terms that best serve its own interest. Later, the same change in economic opportunities can affect the government’s incentive to establish authority inside the region.

2.2 Production and the conservation objective

Conservation is defined by physical activity at the protected sites.

Assumption 2.4 (A specified physical conservation target). Let $\mathcal{E}$ be the designated set of ecologically critical Selva sites. Each site has a stated safe physical allocation set

$\mathcal{A}_r^{\text{safe}}$, describing extraction, land use, inputs and livelihood requirements. Conservation means that the reconstructed allocation belongs to this set at every designated site under the asserted date or limiting criterion. Fiscal gains, certified labels and a lower global contamination fraction are not themselves this condition.

The final result gives finite dates for reaching each positive tolerance of aggregate disturbance and a separate condition on cumulative damage. With imperfect intervention, some ungoverned sites can remain at every finite date. Safe activity also requires the stated production locations and funded household expenditures.

The physical conservation target in Assumption 2.4 makes the allocation inside this boundary relevant to global conservation. Output, mining labor, purchased inputs and operating sites describe different physical margins of damage. Saving a harmful input can improve environmental outcomes at unchanged output; opening a site can increase damage without changing measured productivity. We therefore keep the household and site allocation alongside the quantity of mineral.

The regional problem is to determine which of these allocations the gate can induce, and at what cost, for a given downstream return. This connects an exterior instrument to the activities on which conservation depends.

2.3 Households, production and the two arms of the gate

The gate first chooses which villages to serve from a finite potential set $\mathcal{V}_J$. Its active network is $A \subseteq \mathcal{V}_J$. Excluded villages retain their specified outside opportunities and neither supply nor trade with this network. Coverage is an economic choice because specialization can make villages complementary: two villages may survive through mutual trade even when neither can survive alone.

At village $v \in A$, household utility is

$$u_v = \log(C_v - c_v) + \eta_v \log r_v + \beta_v \log(b_v + b_{0v}), \quad C_v > c_v > 0, \quad r_v > 0, \quad b_v \geq 0 \quad (1)$$

The positive weights η_v and β_v govern the value of leisure and retained cash. The basic consumption basket requires $(1 - \alpha)C_v$ units of a local good and αC_v units of imported necessities, where $0 < \alpha < 1$. Households buy the least-cost bundle producing this Leontief basket, so α measures a physical requirement. Leisure is r_v . Cash b_v buys outside consumption at fixed terminal purchasing power. Its positive utility allowance b_{0v} cannot finance current spending; opening wealth is $a_v \geq 0$.

The household owns its productive assets. It can produce the local good using $F_v(l_v^x)$ and mineral using $G_v(l_v^m, g_v; K_v, A_D)$, subject to

$$l_v^x + l_v^m + r_v \leq 1 \quad (2)$$

Continuous, concave technologies combine household labor with inherited equipment K_v and the dirty technological frontier A_D . Operating this equipment uses the current con-

sumable input g_v , such as fuel or maintenance. The household owns the productive assets; the later research problem specifies separately who can appropriate the returns to an innovation.

The group posts a common mineral procurement price $\theta_B > 0$ and controls both import margins:

$$p_v^z = c_v^z + \mu_z > 0, \quad p_v^g = c_v^g + \mu_g > 0, \quad \mu_z \geq 0 \quad (3)$$

Positive delivered resource costs c_v^z, c_v^g can differ across villages. The productive-input margin can be negative while the delivered price remains positive. Such a subsidy can be privately profitable: the gate recovers part of the benefit when it purchases the additional mineral. Procurement and import prices must therefore be chosen together.

Let p_v be the local-good price and $\bar{p}_v = (1 - \alpha)p_v + \alpha p_v^z$ the basket price. The household-owner solves

$$\max u_v \quad \text{subject to} \quad \bar{p}_v C_v + p_v^g g_v + b_v \leq a_v + \theta_B G_v(l_v^m, g_v) + p_v F_v(l_v^x) \quad (4)$$

together with nonnegativity, time and technological feasibility. The household finances expenditure from opening wealth and gross production income. Decreasing returns leave a return to its fixed productive assets as well as to labor. Both are available for consumption and both affect its supply of mineral.

Proposition 2.5 (Household choice and competitive production). *At positive prices, suppose the relevant feasible choices are bounded and there is a strictly surviving feasible plan. The household problem attains a maximum and has a unique consumption–leisure–cash triple. With differentiable concave technologies and a constraint qualification, its necessary and sufficient conditions are*

$$\begin{aligned} \lambda_v \bar{p}_v &= (C_v - c_v)^{-1}, & \zeta_v &= \eta_v / r_v, \\ \lambda_v p_v F_v' &\leq \zeta_v, & \lambda_v \theta_B G_{vL} &\leq \zeta_v, \\ \theta_B G_{vg} &\leq p_v^g, & \lambda_v &\geq \beta_v / (b_v + b_{0v}) \end{aligned} \quad (5)$$

Each activity inequality binds when the activity is positive, with the appropriate normal-cone condition at a technological bound. The budget binds. Competitive firms paying $w_v = \zeta_v / \lambda_v$ and returning their profits to the household decentralize these choices.

Proof. The feasible budget set is convex because expenditure is linear and production income is concave. Extend utility to minus infinity at the essential boundaries $C_v = c_v$ and $r_v = 0$. A maximizing sequence cannot approach either boundary while its other arguments are bounded above and a strictly surviving plan is available. A convergent subsequence therefore attains the maximum in the utility domain. Strict concavity in (C_v, r_v, b_v) makes that triple unique. Since additional cash is valuable, no budget slack is optimal. Differentiation yields (5); the supporting inequalities of the concave program establish sufficiency.

The local firm maximizes $p_v F_v(l) - w_v l$ and the mining firm maximizes $\theta_B G_v(l, g) - w_v l - p_v^g g$. Giving both profits to their owner and substituting labor supply $1 - r_v = l_v^x + l_v^m$ into the household's wage-plus-dividend budget recovers (4). The same supporting conditions make the firms' factor choices optimal. Wages and dividends are an alternative decomposition of production income, not an addition to it. $\square$

The household conditions show why a procurement-price change has no independent interpretation. The necessity price determines what the mining return buys, while the input price affects the resources required to earn it. Local prices and participation complete this system. Regional supply is the outcome of all these terms, rather than a household response to θ_B holding the other terms implicit.

2.4 Internal trade, participation and costly coercion

Local goods move on a specified finite directed village network. Shipping t_{vw} units from v delivers t_{vw}/δ_{vw} to w , with $\delta_{vw} \geq 1$. A competitive linked equilibrium combines household optima with market clearing and zero-profit shipment conditions:

$$\begin{aligned} F_v + \sum_w t_{wv}/\delta_{wv} - \sum_w t_{vw} &= (1 - \alpha)C_v, \\ p_v - p_w/\delta_{vw} &\geq 0, \quad t_{vw}(p_v - p_w/\delta_{vw}) = 0 \end{aligned} \quad (6)$$

The prospect of shipment limits price differences even on routes that carry no trade. Free disposal can be included with zero-value complementarity. When transport is costless, removing inessential shipment cycles preserves household allocations and market-clearing values.

The outside utility available to household v is $\bar{U}_v$. The gate may apply coercive effort e_v at real cost $\Psi_v(e_v)$, so participation requires

$$u_v \geq \bar{U}_v - \Delta_v(e_v), \quad 0 \leq e_v \leq \bar{e}_v \quad (7)$$

The continuous, strictly increasing function Δ_v , with $\Delta_v(0) = 0$, measures the deterioration of the outside option. Its cost Ψ_v is continuous and increasing, with $\Psi_v(0) = 0$. Force can make refusal less attractive, but induced production must still satisfy subsistence, available time and the household's demand for positive leisure.

Proposition 2.6 (Minimum coercion). *If coercion affects the household only through (7), an implementable allocation uses*

$$e_v^* = \Delta_v^{-1}((\bar{U}_v - u_v)_+) \quad (8)$$

whenever this lies in $[0, \bar{e}_v]$. An allocation requiring more than $\bar{e}_v$ is infeasible. At an uncapped, binding participation constraint and on a differentiable allocation branch, $de_v^/ds = -u_{v,s}/\Delta_v'(e_v^*)$ for a regional control s .*

Proof. The feasible efforts form an upper interval whose smallest element is (8). Increasing effort above it leaves the induced household allocation unchanged and increases its resource cost. The least element is therefore optimal. Implicit differentiation of the binding constraint gives the derivative. At a cap or a slack constraint, the relevant boundary condition replaces that derivative. $\square$

Extracting a larger trading margin can therefore raise the cost of maintaining control. Worse household terms require greater coercion when participation binds, and the effort cap can make continued participation infeasible. The gate compares village networks by total profit, including the terms and coercion needed to keep their households active.

2.5 Subsistence, distress and trade-dependent survival

The importance of subsistence can be seen before adding the adjustments in local trade and purchased inputs. The connection is to the household survival motive emphasized by Basu and Van [1998]; the present question concerns the household's mining effort. Temporarily suppress local trade, productive inputs and the utility of cash. Let mining be linear, with earning power $w = \theta_B A$, and let producing the local complement of one basket take $d = (1 - \alpha)/A_x$ units of time. Write $B = \alpha p^z$ for the imported component's price. With wealth a , positive mining labor l finances $BC = a + wl$, while $r = 1 - dC - l$. Define $R = 1 - da/B$ and $k_1 = 1 + dw/B$ on a strictly feasible interior branch.

Proposition 2.7 (The subsistence labor response). *In this benchmark the interior optimum satisfies*

$$l = \frac{R}{(1 + \eta)k_1} + \frac{\eta(Bc - a)}{(1 + \eta)w}, \quad l_w = -\frac{Rd}{(1 + \eta)Bk_1^2} - \frac{\eta(Bc - a)}{(1 + \eta)w^2} \quad (9)$$

In particular, if $a \leq Bc$, a higher mining return reduces required mining labor on that branch. More generally, the response is positive exactly when $a - Bc > Rdw^2/(\eta Bk_1^2)$, negative when the inequality is reversed, and zero at equality.

Proof. Substitute $C = (a + wl)/B$ and $r = R - k_1 l$ into utility. The first-order condition is $w/(a + wl - Bc) = \eta k_1/(R - k_1 l)$. Solving this linear equation in l gives the first expression; differentiating while holding a, B, d fixed gives the second. Its sign gives the threshold. Strict concavity validates the interior first-order condition. $\square$

When other resources fail to cover essential imports, both terms in the displayed labor derivative are weakly negative. A higher mining return allows the household to secure its necessities with less work. Once wealth covers subsistence, the substitution effect can instead dominate. This household response is one component of the gate's decision: the gate also accounts for import margins, participation and the value of the mineral it obtains.

In the linked economy, common terms also change local prices, trade, input use, retained cash and the coercion needed to secure participation. Collect the household and market-clearing conditions on a regular branch as $\mathcal{F}(x; s) = 0$. If $\mathcal{F}_x$ is nonsingular,

$$x_s = -\mathcal{F}_x^{-1}\mathcal{F}_s, \quad q_s = \sum_v (G_{vL}l_{v,s}^m + G_{vg}g_{v,s}) \quad (10)$$

This total derivative defines normal, distress and flat responses in the linked economy. Local prices p_v adjust with the household choices they support. At village entry or a change in binding constraints, the gate compares the resulting feasible allocations directly.

Trade can determine whether survival is feasible at all. Without outside wealth, a farming village needs mineral earnings to finance imported necessities; a mining village without local production needs the farmer's good. Trading together allows the miner's export earnings to finance both households' requirements. Appendix A constructs a smooth two-village economy with a unique market-clearing local price. The economically viable unit is then the village network, rather than either village in isolation.

2.6 The regional resource account

An allocation delivers

$$q = \sum_{v \in A} G_v, \quad Z = \sum_{v \in A} \alpha C_v, \quad I = \sum_{v \in A} g_v \quad (11)$$

The gate pays $\theta_B q$ for mineral, earns the necessity and input margins, and pays mineral transport, coverage and coercion costs. Let z collect its terms, active set, coercion and the induced competitive equilibrium, and let $T(z) \geq 0$ be its transport and coverage bill. The cost of implementing this allocation is

$$C(z) = \theta_B q - \mu_z Z - \mu_g I + T(z) + \sum_v \Psi_v(e_v) \quad (12)$$

Import margins return part of the mineral payment to the gate. Its net implementation cost is therefore smaller than gross procurement expenditure when those margins are positive. Consolidating the regional budgets shows which expenditures survive the cancellation of internal transfers.

Proposition 2.8 (Consolidated resources). *Every linked equilibrium satisfies*

$$\theta_B q + \sum_v a_v = \sum_v (p_v^z \alpha C_v + p_v^g g_v + b_v), \quad (13)$$

$$C(z) = \sum_v (c_v^z \alpha C_v + c_v^g g_v + b_v - a_v) + T(z) + \sum_v \Psi_v(e_v) \quad (14)$$

The identities hold with iceberg losses on local routes.

Proof. Sum the binding household budgets. The uncanceled local-good value is

$$\sum_v p_v [F_v - (1 - \alpha)C_v] = \sum_{vw} t_{vw}(p_v - p_w/\delta_{vw}) = 0$$

where the first equality uses physical clearing and the second uses route complementarity. This proves (13). Substituting $p_v^z = c_v^z + \mu_z$ and $p_v^g = c_v^g + \mu_g$ into (12) proves (14). The physical transport loss $\sum_{vw}(1 - 1/\delta_{vw})t_{vw}$ remains real even though its counterpart in local resident payments cancels. $\square$

The account makes the distributional role of prices explicit. Tied imports can absorb a high nominal mineral payment, whereas a lower input price can reduce the resources needed to induce output. Consumption, cash, inputs and coercion in (14) reveal what these terms implement. Cash retained by households enters the gate's private cost alongside real resource expenditure; the expression is not a measure of regional welfare.

If $W(q)$ is the gross value of delivering stock q to the downstream stage, the group's objective is $\Pi(z) = W(q(z)) - C(z)$. On a smooth regime its derivative in a control s_j is

$$\begin{aligned} \Pi_j = & (W_q - \theta_B)q_j - \mathbf{1}_{\{s_j=\theta_B\}}q + \mathbf{1}_{\{s_j=\mu_z\}}Z + \mu_z Z_j \\ & + \mathbf{1}_{\{s_j=\mu_g\}}I + \mu_g I_j - T_j - \sum_v K'_v(e_v)e_{v,j} \end{aligned} \quad (15)$$

The gate balances the marginal downstream return against procurement, both import activities, and the cost of maintaining participation. A binding input-price restriction or participation cap changes that balance; a coverage change requires comparing the total profits of the candidate networks.

2.7 The cost of implementing supply

To choose output, the gate must be able to implement the associated local allocation. Let $\mathcal{Z}$ collect the competitive equilibria supported by its terms. Where an institution selects among multiple equilibria, its selection rule is fixed independently of the downstream policy. This keeps the opportunities for regional implementation constant in the policy comparison.

For output q with at least one admissible implementation, define

$$\mathcal{C}(q) = \min_{z \in \mathcal{Z}: q(z)=q} C(z), \quad \Gamma(q) = \arg \min_{z \in \mathcal{Z}: q(z)=q} C(z) \quad (16)$$

The minimum in this expression must correspond to an allocation the gate can attain. Subsistence, resource costs and diminishing returns provide economic bounds on the choices in the following smooth class. A closed equilibrium graph then preserves feasibility at their limits.

Assumption 2.9 (A smooth linked production class). The active network is finite and strongly connected with finite iceberg factors. Production is $F_v(l) = A_v^x l^{\nu_v}$ and $G_v(l, g) =$

$A_v^m l^{\gamma_v} g^{\kappa_v}$, where $0 < \nu_v < 1$, $\gamma_v, \kappa_v > 0$ and $\gamma_v + \kappa_v < 1$. Zero coefficients permit specialization; at least one local producer and one miner are capable. A village with the lowest delivered input resource cost has positive mining capability. Wealth is finite, coercion is bounded, and

$$d_A := \sum_{v \in A} \alpha c_v^z c_v - \sum_{v \in A} a_v > 0 \quad (17)$$

Transport and coercion costs are continuous and nonnegative. The equilibrium graph is the full competitive correspondence or an independently fixed selection whose graph is closed on the relevant sublevels.

Theorem 2.10 (Attained regional cost). *Under Assumption 2.9, every nonempty implementation-cost sublevel with $q \geq q_0 > 0$ is compact in physical allocations and prices, after inessential shipment witnesses are projected out. Every nonempty positive-output fiber attains its minimum. If $W(q) \leq \bar{v}q$ for finite $\bar{v} > 0$, every operating choice that can beat the zero-profit exit option has an economic output floor $q \geq d_A/\bar{v}$ and belongs to one bounded cost sublevel. Upper semicontinuity of the actual downstream continuation then gives an attained profit maximum, including finite coverage comparisons among branches satisfying these conditions and the separate exit option.*

Economic argument. The resource identity bounds consumption, purchased inputs and cash on a cost sublevel. Bounded coercion gives a finite lower participation utility. Together these bounds separate consumption and leisure from their singular subsistence boundaries. Valuable cash then bounds shadow wages above, while the imported necessity floor bounds them below. Decreasing-returns mining and its factor conditions bound the common procurement quote on a positive- output fiber; they also bound delivered input prices. Necessity margins are bounded by the external cash account. Local production and the connected trade inequalities bound local prices.

The resulting household, participation and market-clearing conditions have a closed graph inside those economic bounds. Shipment cycles can be removed without changing the allocation or resource bill. A minimizing sequence has a convergent implementable limit. For nonloss operation, $C(z) \geq d_A$ and $W(q) \leq \bar{v}q$ supply the output floor. Sublinear input productivity ensures that the linear resource bill eventually exceeds the bounded-value output it can produce, giving one upper cost bound for all choices that can beat exit. The detailed inequalities, including cash and trade corners, are derived in Appendix A. $\square$

An operating region must earn enough outside income to finance the subsistence deficit in (17). This requirement gives the output floor its economic meaning, including when farmers and miners finance it jointly through trade. The theorem establishes attainment on each nonempty implementation graph. The appendix's linked-price construction supplies such a graph, with a continuous selected equilibrium, for a smooth subclass.

The resulting allocation need not identify each gross quote. Suppose that a village uses $G(l, g) = A \min\{l, g\}$, so efficient use has $g = l$. Its net mining earning power is $w = A\theta_B - p^g$. Raising the procurement quote by t and the input quote by At leaves w

unchanged. Both the household problem and the resource cost are identical. The same allocation can therefore be supported by several gross price pairs, even with positive retained cash. Strict decreasing returns removes this unbounded offset ray, while still permitting multiple bounded equilibrium prices. The economically identified objects are the allocation and the set of terms supporting it; uniqueness of a nominal quote requires its own argument.

2.8 The mineral delivered to the world market

The regional problem supplies a cost and an allocation for each attainable delivery. The set $\Gamma(q)$ records the terms, household choices, coverage and coercion that attain the lowest net cost. After procurement, the seller values the physical stock through its available trading institutions. If it can commit to offered capacity n , write

$$W(q) = \max_{n \in \mathcal{N}(q)} W_n(q), \quad \mathcal{N}(q) \subseteq \{0, \dots, \lfloor q/\bar{q} \rfloor\} \quad (18)$$

The allowable capacities depend only on physical output. Offering all complete packages gives $\mathcal{N}(q) = \{\lfloor q/\bar{q} \rfloor\}$; allowing withholding enlarges this pre-bid choice set.

Theorem 2.11 (Regional output as an implemented choice). *Suppose the output fibers and downstream capacity maxima attain and the regional maximum exists under the conditions above. The full joint choice of posted terms, village coverage, coercion, equilibrium allocation and offered capacity is equivalent to*

$$q^* \in \arg \max_{q \in \mathcal{Q} \cup \{0\}} \{W(q) - \mathcal{C}(q)\}, \quad z^* \in \Gamma(q^*) \quad (19)$$


with the separately specified exit payoff at $q = 0$. An unchanged optimal quantity leaves the implementation correspondence $\Gamma(q^)$ unchanged in a pure downstream experiment. It implies unique unchanged posted terms only when those components of $\Gamma(q^*)$ are single-valued.*

Proof. Every implementation of q costs at least $\mathcal{C}(q)$ and every allowed capacity produces at most $W(q)$. Its profit is therefore bounded above by the displayed objective. An attained cost-minimizing implementation and an attained capacity choice achieve that bound at each q . Maximizing over outputs and comparing exit proves equivalence. The fiber problem contains only the unchanged regional primitives, so its solution set at a fixed output is unchanged. Set invariance does not imply a single selected quote. $\square$

The same choice applies when the downstream institution is a committed continuum menu with an attained resale equilibrium $W^M(q, \tau)$. Its compact price-and-allocation construction supplies upper semicontinuity. Bounded feasible output, or cost that dominates large-stock continuation, then yields a regional maximum. The resale example uses the household-owned input technology derived in the power-envelope proof and includes the household's operating rent in implementation cost.

The supply of mineral thus carries with it a feasible pattern of livelihoods and control. A change in $W(q)$ affects the gate's ranking of attainable quantities. At an unchanged quantity, the correspondence of cost-minimizing implementations remains the same; a unique pattern requires a unique implementation or an unchanged selection from that set. We next derive the market opportunities behind this value and ask when a selective restriction on buyers changes the return to extraction.

3 Selective regulation and the return to extraction



Retention and competing buyers shape how selective regulation reaches the gate.

3.1 Alternative buyers and alternative uses of stock

A restriction on some buyers leaves the intermediary with a choice. It can sell to other buyers, conceal the mineral's origin, or retain stock for a later opportunity. Each alternative affects the return to producing the mineral in the first place. The effect of selective regulation therefore depends on the market opportunities that continue to support extraction.

The spatial price spread $0 < \theta_B < \theta_E$ creates the initial trading opportunity. Mineral purchased inside the region at θ_B can command the outside quote θ_E once it becomes an equivalent standardized deliverable. In the liquid-market benchmark, changing the owner of an eligible delivery causes no further discount. The regional lot is small relative to the market, and the quote is fixed during the trading round. World prices can still adjust across dates and policies.

The depth of the gold market motivates this assumption. LBMA describes trading between market institutions once eligible Good Delivery bars enter an approved London vault [London Bullion Market Association, 2026]. Reaching that market still requires refining, delivery, finance, compliance and access, whose costs enter the model. Other

minerals have different grades and trading institutions. The LME's cobalt liquidity window and the IEA's account of mineral markets illustrate these differences [[London Metal Exchange, 2026](#), [International Energy Agency, 2025](#)].

Two forces can protect the return to illegal production. The value of retention supports inventory and sets a floor under auction payments. Competition among unaffected buyers can sustain those payments even when laundering can no longer expand. Whether these forces preserve extraction depends on a further comparison: how much the policy lowers profit at the chosen inventory relative to every alternative inventory.

Initially, the world resale quote θ_E , laundering technology and regional primitives are fixed. A fixed set of regulated buyers $\mathcal{R}$ faces compliance cost $\tau \geq 0$ per standardized unit acquired; buyers in $\mathcal{U}$ remain unaffected. This is the policy available under the initial restriction on enforcement reach. Certification will subsequently change feasible resale and, with it, market equilibrium.

Physical stock is divided into complete packages and a possible residual:

$$Q = \lfloor q/\bar{q} \rfloor, \quad r_q = q - \bar{q}Q, \quad 0 \leq r_q < \bar{q} \quad (20)$$

Quotes θ_B and θ_E are prices per standardized physical unit; marginal buyer values are amounts per package, and transfers are total payments. Procurement expenditure $\theta_B q$ is paid before bidding and enters the regional implementation cost once.

Buyers are finitely many distinct owners, collected in $\mathcal{I}$, with $N = |\mathcal{I}|$. Buyer i has public capacity $\bar{k}_i$ and private type ω_i . Its operating value $J_i^0(x; \omega_i)$ accounts for processing, finance and resale before payment for the auctioned mineral. The corresponding marginal package values are

$$\begin{aligned} b_i^k(\omega_i) &= J_i^0(k\bar{q}; \omega_i) - J_i^0((k-1)\bar{q}; \omega_i), \\ v_i(k; \omega, \tau) &= \sum_{j=1}^k b_i^j - \bar{q}\tau k \mathbf{1}_{\{i \in \mathcal{R}\}} \end{aligned} \quad (21)$$

Concavity of J_i^0 makes these increments decrease. A buyer whose post-policy value becomes negative can receive zero units. Access to an unconstrained outside source of identical mineral at θ_E bounds its opportunity-adjusted marginal package value by $\bar{q}\theta_E$. The value of access inside the auction depends on access outside it.

All inventories and policy levels use the same buyers, capacities and type law ν . We impose the integrability condition

$$\mathbb{E}_\nu \left[\sum_{i \in \mathcal{I}} \sum_{k=1}^{\bar{k}_i} (b_i^k)_+ \right] < \infty$$

Finite retention value and this integrability condition make expected receipts and their policy changes finite. Procurement uses the information available before bidding. A conditional analysis can set $\nu = \mu(\cdot \mid A_0)$ for an initially specified positive-probability

event A_0 . Keeping that event fixed across inventories and policies ensures that the same buyer population and information govern every production comparison.

3.2 What the seller can earn by retaining mineral

The relevant alternative to an auction sale is the best use of retained stock. The intermediary allocates that stock among laundering, storage and disposal, obtaining gross value

$$H(y) = \max_{\ell+s+d=y} \{\theta_E \ell - V(\ell) + S(s)\}, \quad \ell, s, d \geq 0 \quad (22)$$

Laundering quantity ℓ incurs real cost $V(\ell)$, storage s has net continuation value $S(s)$, and disposal d is free. Channel capacities constrain the same allocation problem. We assume that H is finite, continuous, concave and nondecreasing on the relevant domain. Convex laundering costs, concave storage returns, convex capacities and free disposal produce these properties when the optimum is attained.

Unsold packages and residual stock r_q compete for the same channels. The value of returning a package to the seller consequently depends on the stock it already retains. Procurement is sunk at this stage, so the package's reservation value is its incremental gross retention return.

This return is $H(y + \bar{q}) - H(y)$. When laundering alone is chosen at both endpoints, it equals $\bar{q}\theta_E - [V(y + \bar{q}) - V(y)]$. It exceeds the package procurement bill precisely when incremental laundering cost is below $\bar{q}(\theta_E - \theta_B)$ (Lemma B.1). The spatial spread therefore creates a potential market outlet; channel costs determine how much of that spread the seller can preserve.

The seller commits before reports to offered capacity $n \leq Q$, its genuine retention schedule H , and efficient allocation with Clarke pivot transfers. Offering all stock fixes $n = Q$; permitted withholding allows the seller to choose n before bidding. The welfare and receipt functions are

$$S_n(q, \omega, \tau) = \max_{\substack{0 \leq k_i \leq \bar{k}_i, k_i \in \mathbb{Z} \\ \sum_i k_i \leq n}} \left\{ \sum_i v_i(k_i; \omega, \tau) + H(q - \bar{q} \sum_i k_i) \right\}, \quad (23)$$

$$\begin{aligned} p_i &= S_{n,-i} - S_n + v_i(k_i^*; \omega, \tau), \\ R_n(q, \omega, \tau) &= \sum_i p_i + H(q - \bar{q} \sum_i k_i^*), \quad W_n(q, \tau) = \mathbb{E}_\nu R_n(q, \omega, \tau) \end{aligned} \quad (24)$$

The counterfactual $S_{n,-i}$ removes buyer i and all its orders while preserving stock, capacity and retention opportunities. Seller receipts consist of buyer transfers and the value of actual retained stock.

Proposition 3.1 (Incentives within the committed mechanism). *Truthful reporting is weakly dominant for every buyer, and a truthful buyer has nonnegative utility. These properties hold after the announced common compliance cost. They do not imply that the mechanism is the unique revenue-maximizing institution available to the seller.*

Proof. Fix other reports, q, n, H and τ . Substituting the pivot transfer into the buyer's true utility gives the total true welfare of the resulting allocation minus $S_{n,-i}$. The subtracted quantity does not depend on the buyer's report. Truthful reporting maximizes the first quantity over the unchanged feasible set, proving dominant-strategy truthfulness. The allocation that excludes the buyer remains feasible with that buyer present, so truthful utility $S_n - S_{n,-i}$ is nonnegative. These are statements about the committed allocation and transfer rule; neither compares the seller's receipts with every other individually rational mechanism. $\square$

3.3 When genuine-retention VCG is an optimal selling institution

Valuable retention explains the seller's reservation orders. Optimality of the selling institution depends in addition on the information rents that buyers must receive. Following the complete-game approach of [Ausubel and Cramton \[1999\]](#) and the efficient-payment comparison of [Krishna and Perry \[1998\]](#), we compare procedures whose complete outcomes allocate stock efficiently. The seller has public nonlinear retention value H , and buyers pay the whole-buyer transfers in (24).

Assumption 3.2 (Fixed-stock efficient-resale benchmark). Initially $\tau = 0$ and q is already procured. The following restrictions apply to the institution comparison.

1. True H , stock, capacities and the buyer laws are public before mechanism choice. Buyer marginal vectors are independent private types on compact convex domains of nonnegative decreasing values containing zero. All types, including zero, can abstain from every trade at zero utility.
2. The seller commits to initial terms and cash settlement. All Q packages, including initially unoffered or unsold packages, remain available for voluntary terminal trade before any irreversible use. Final feasible quantities and channel opportunities are the same across procedures.
3. Private transfers balance; preferences are quasilinear; all costs and outside operating opportunities are counted once. Initial nonparticipants can join resale. There is no new outside signal revealing true types after auction reports.
4. The comparison is over complete-game Bayesian Nash equilibria whose terminal allocations are efficient relative to true values and H , as in Definition B.2. This is a substantive perfect-resale outcome restriction, not a claim that liquidity or a finite bargaining protocol guarantees efficiency after every selling procedure.

Theorem 3.3 (Optimality of genuine-retention VCG under efficient resale). *Under Assumption 3.2, every admissible complete outcome gives the seller expected total continuation payoff at most $W_Q(q, 0)$. The all-offered VCG auction with genuine H and the whole-buyer pivot payments in (24) attains this bound. Specifically, when terminal trade*

requires final voluntary consent to a balanced contract against initial ownership, truthful buyer reports and no trade form an ex post Nash equilibrium of the fixed-procedure auction-and-resale game. The equilibrium resists combined misreport/resale and non-participation/resale deviations. No additional strategic reserve premium or participation charge is needed on this domain.

For the announcement comparison, write the actual complete-game payoff as

$$\Pi_q(\widehat{H}; H) = \mathbb{E} \left[T_0^{\text{net}}(\widehat{H}) + H(y^{\text{final}}(\widehat{H})) \right]. \quad (25)$$

The net transfer T_0^{net} includes seller payments and receipts throughout auction and subsequent trade, evaluated at the induced complete equilibrium. Actual retained stock is valued at true H , the second argument of the payoff function.

Corollary 3.4 (An optimal genuine announcement). *Under Assumption 3.2, select the attaining equilibrium at the genuine announcement. For every alternative announcement and each of its admissible complete equilibria,*

$$\Pi_q(H; H) = W_Q(q, 0) \geq \Pi_q(\widehat{H}; H). \quad (26)$$

Thus declaring the joint retention increments without an additional strategic premium is an expected-payoff optimal choice. Neither uniqueness nor dominant-strategy truthfulness of the revenue-receiving seller is asserted.

The economic force behind the bound is the participation rent of the lowest type. On the stated domain, every efficient incentive-compatible procedure gives each buyer interim expected utility equal to its interim pivot utility plus a nonnegative participation constant. Balanced cash transfers make any such constant a deduction from the seller's payoff. VCG sets the constants to zero. Appendix B.1 establishes this payment bound and shows that truthful reports followed by no trade attain it, including against joint reporting-and-trading deviations. The equilibrium claim concerns ex post Nash behavior in the stated procedure; off-path sequential implementation and the ranking of inefficient bargaining outcomes require additional structure.

Public retention and independent buyer identities perform distinct roles. Independently disclosed H fixes the seller's opportunity cost before mechanism selection. A privately informed seller instead needs an information rent of its own. The one-buyer, independent-uniform example in Proposition B.7 shows that balanced transfers cannot generally finance both parties' truthful participation rents while preserving efficiency. Coordinated buyers create another difficulty: they can suppress the rival orders that determine payments. The two-buyer example in Appendix B.9 demonstrates this collusion incentive, which individual truthfulness does not resolve.

Proposition B.8 carries the fixed-stock result back to procurement when the bound and its attainment hold at every feasible stock under common information. It also covers initial capacity choices when withheld stock remains available in an efficient continuation.

A selective cost that attaches once to final use and cannot be avoided in resale preserves this comparison. A charge on direct auction acquisitions has different incidence. The following resilience analysis takes the committed auction as given and derives its response under the fixed type law; comparing it with other institutions requires the assumptions above.

3.4 Who owns the demand that sets the price?

For retained stock $y \geq 0$, define the marginal retention value of its j th additional package by

$$\rho_j(y) = H(y + j\bar{q}) - H(y + (j - 1)\bar{q}), \quad j = 1, 2, \dots \quad (27)$$

Let $a = q - n\bar{q}$ denote the stock outside the offered capacity, including the fractional residual. The n optional seller orders are $\rho_1(a), \dots, \rho_n(a)$; the sequence is empty when $n = 0$. For $n \geq 1$, the last and smallest seller order is the one-package reserve $\rho(q) := H(q) - H(q - \bar{q}) = \rho_n(a) = \rho_1(q - \bar{q})$. Concave retention makes these orders nonnegative and decreasing. Combining them with buyer orders and selecting the highest n implements efficient allocation, with fixed priority at ties and prefixes preserved within each agent. The constant $H(a)$ values unavoidable retention. Selected seller orders add the optional packages it keeps.

Lemma 3.5 (Whole-buyer critical payments). *For buyer i receiving $k_i^* > 0$ packages, let $z_{-i,[j]}$ be the j th highest order after removing all of i 's orders, including its losing ones. Index these packages by $u = 1, \dots, k_i^*$. Its transfer is*

$$p_i = \sum_{u=1}^{k_i^*} c_{iu}, \quad c_{iu} = z_{-i,[n-k_i^*+u]} \quad (28)$$

The terms in this sum are the additional rival orders filled when buyer i is removed.

Proof. Telescoping the seller orders gives $H(a + j\bar{q}) - H(a)$, so maximizing (23) is equivalent to selecting the highest n decreasing orders. The selected rival orders other than i occupy the first $n - k_i^*$ positions in its rival book. Removing i fills the next k_i^* positions. If B_i is the sum of the previously filled rival orders and N_i the sum of those newly filled, then $S_n = H(a) + v_i(k_i^*) + B_i$ and $S_{n,-i} = H(a) + B_i + N_i$. Their pivot difference is exactly N_i , proving the formula. $\square$

The ownership of demand determines the seller's competitive alternatives. A single buyer pays nothing when retention is worthless, however strong its own excess demand: its losing orders disappear along with its winning orders in the exclusion comparison. The same marginal values held by independent buyers can support positive revenue. For resilience, the relevant depth is therefore the demand surviving the removal of an owner.

Proposition 3.6 (Reallocation toward unaffected buyers). *At fixed q, n, H and realized types, an unaffected buyer's allocation weakly increases when τ increases, with the same*

fixed priority at ties. Under any fixed type law its post-policy allocation therefore first-order stochastically dominates its pre-policy allocation.

Proof. Every regulated marginal order weakly decreases while the unaffected buyer's own orders, the other unaffected orders and retention orders stay fixed. In the top- n representation this can only improve the position of each of that buyer's orders relative to competing orders. Every own order selected before the change is still selected afterward. This gives the statewise inequality, which implies each corresponding tail-probability inequality under the fixed type law. $\square$

3.5 Retention as a floor under payments

Let $m = \sum_i k_i^*$ and $y = q - \bar{q}m$ be actual external sales and retained stock. For the units assigned to a winning buyer define

$$D_i(y) = H(y + k_i^* \bar{q}) - H(y) = \sum_{u=1}^{k_i^*} \rho_u(y) \quad (29)$$

This block opportunity value measures what the seller would gain by receiving the buyer's mineral back. When laundering is the sole interior retention choice, the unit opportunity price is

$$\frac{\rho_u(y)}{\bar{q}} = \theta_E - \frac{V(y + u\bar{q}) - V(y + (u-1)\bar{q})}{\bar{q}}$$

As retained stock increases, rising laundering costs can reduce the reservation price even while the channel remains profitable. The world quote determines the gross opportunity; the cost of using it determines the marginal reserve.

Proposition 3.7 (The opportunity floor). *Every winning buyer satisfies*

$$c_{iu} \geq \rho_u(y), \quad p_i \geq D_i(y), \quad R_n(q, \omega, \tau) \geq H(q) \quad (30)$$

Thus a positive retention opportunity margin above acquisition cost supports a positive acquisition margin on the externally sold block. A strictly positive auction premium over retention requires strictly stronger rival orders; a winner's own willingness to pay above the reserve does not suffice.

Proof. The unselected seller orders available when winner i is removed include the consecutive returns $\rho_1(y), \dots, \rho_{k_i^*}(y)$. The newly filled rival orders are the highest available orders, so their respective values are at least these returns. Summing proves the block inequality. For total stock, set $d_i = k_i^* \bar{q}$. Concavity gives decreasing increments, and therefore

$$H(q) - H(y) \leq \sum_i [H(y + d_i) - H(y)] = \sum_i D_i(y) \leq \sum_i p_i$$

Adding $H(y)$ proves the receipt floor. An allocation in which the returned seller order itself sets the critical payment attains equality, regardless of how high the winning value is. $\square$

The acquisition margin on a sold block has two components:

$$\frac{p_i}{k_i^* \bar{q}} - \theta_B = \underbrace{\frac{D_i(y)}{k_i^* \bar{q}} - \theta_B}_{\text{retention acquisition margin}} + \underbrace{\frac{p_i - D_i(y)}{k_i^* \bar{q}}}_{\text{auction premium}} \quad (31)$$

Retention can sustain a positive acquisition margin even when the weakly nonnegative auction premium is zero. Whether that margin supports production depends on the full regional cost, including import receipts, transport and coercion. The payment calculation identifies the value of mineral; Selva's implementation problem determines the cost of obtaining it.

Buyer-specific opportunity costs also differ from the value of retaining all stock at once. Each exclusion returns one buyer's block to the same initial stock y . Concavity allows the sum of these separate returns to exceed the joint return from all sold units. The aggregate premium $R_n - H(q)$ can thus be positive even when each buyer pays exactly its block opportunity value. The appendix decomposes this difference.

3.6 Competition after laundering stops expanding

Let $\ell(y)$ denote the selected laundering quantity in (22). Use $p^R = \theta_E - \tau$ as a scalar index of regulated buyers' unit values. This index is not a common auction payment. Actual laundering is

$$L(q, p^R, \omega) = \ell(q - \bar{q}m(q, p^R, \omega)) \quad (32)$$

At fixed stock, laundering is locally constant whenever the integer allocation stays unchanged. A binding laundering capacity also limits its response when additional stock is retained or disposed. Buyers can continue to compete for mineral in either case.

Proposition 3.8 (A zero direct laundering response). *On a differentiable regime with fixed offered capacity,*

$$L_{p^R} \Big|_q = -\bar{q}\ell'(y)m_{p^R} \Big|_q, \quad \frac{dL(q^*(p^R), p^R)}{dp^R} = L_{p^R} + L_q \frac{dq^*}{dp^R} \quad (33)$$

Thus a zero direct laundering response is consistent with changing payments, positive profits and changing regional production. Under uncertainty, statewise constancy between allocation switches does not generally imply constancy of expected laundering.

Proof. The two equalities are the partial and total chain rules applied to (32). No profit or payment equation appears in either. Under uncertainty, the type threshold defining an allocation can move with policy: the probability mass switching allocations can change

even if every fixed type has a locally constant integer allocation except at its own switch. Appendix B gives an exact continuous-distribution example. $\square$

Define the unaffected rival depth for $n \geq 1$ by

$$e_{\mathcal{U}}(n, \omega) = \min_i b_{\mathcal{U} \setminus \{i\}, [n]} \quad (34)$$

Here the subscript selects the n th highest pooled marginal order, with value zero if fewer than n orders exist. Adding the seller's n orders to each unaffected rival book gives minimum $e_{\mathcal{U}, H}(q, n, \omega)$. At fixed inventory and retention technology, the regulated-buyer cost changes neither floor.

Proposition 3.9 (The remaining-market floor). *For every externally sold package and every $\tau \geq 0$,*

$$c_{iu} \geq \max\{\rho_u(y), e_{\mathcal{U}, H}(q, n, \omega)\}, \quad e_{\mathcal{U}, H} \geq e_{\mathcal{U}} \quad (35)$$

In particular, if $e_{\mathcal{U}, H} > \bar{q}\theta_B$, every externally sold package has a strictly positive acquisition margin, including a regime on which the direct laundering response is zero.

Proof. The critical rank $n - k_i^* + u$ is at most n . Its rival pool contains all the unaffected orders other than i 's and the complete seller schedule, so its order statistic is at least the n th order of that sub-book. Including retention orders can only raise the unaffected-book floor. Combine this observation with (30) and divide the package price by $\bar{q}$ to compare it with θ_B . $\square$

Competition outside the regulated group can therefore sustain payments after the laundering response has ended. For packages sold with certainty, an expected floor above acquisition cost supports a positive expected placement margin. When the occurrence or number of sales varies by state, the relevant expectation instead averages the floor over packages actually sold. Realized margins can still differ across market states. If n buyer orders also exceed the highest seller order in every state, the full offered inventory is sold externally with certainty.

Positive margins leave the production question open. A smaller inventory can become more profitable even while every package in the original one sells above acquisition cost. We now compare the loss in receipts across inventories, connecting the market's response to the gate's production choice.

3.7 The demand that transmits policy to payments

Let k^{-i} be the allocation that excludes buyer i , with $k_i^{-i} = 0$. Write $m_{\mathcal{R}} = \sum_{j \in \mathcal{R}} k_j^*$ and $m_{\mathcal{R}}^{-i} = \sum_{j \in \mathcal{R}} k_j^{-i}$. Define exposure

$$\chi_n(q, \omega, \tau) = \sum_i m_{\mathcal{R}}^{-i} - (N - 1)m_{\mathcal{R}} = \sum_i \sum_{j \in \mathcal{R} \setminus \{i\}} (k_j^{-i} - k_j^*) \quad (36)$$

Exposure measures the regulated demand that determines critical payments. An order can contribute to several buyers' transfers and is counted in each such comparison. Its weight in payments therefore differs from the physical quantity acquired by regulated buyers.

Theorem 3.10 (The exact exposure identity). *For fixed q, n, ω , the seller's gross receipts are continuous and piecewise affine in τ . Away from their finitely many breakpoints,*

$$\partial_\tau R_n = -\bar{q}\chi_n, \quad 0 \leq \chi_n \leq \sum_i k_i^* \leq n \quad (37)$$

Let $\bar{\chi}_n(q, \tau)$ denote expected exposure of payments to regulation under the fixed information law ν . Then

$$W_n(q, \tau) - W_n(q, 0) = -\bar{q} \int_0^\tau \bar{\chi}_n(q, t) dt, \quad \bar{\chi}_n = \mathbb{E}_\nu \chi_n \quad (38)$$

Proof. Summing the pivot transfers and adding actual retention yields

$$R_n = \sum_i S_{n,-i} - (N-1)S_n \quad (39)$$

Each welfare value is a maximum over finitely many affine functions of τ . On an open segment its derivative is minus $\bar{q}$ times the regulated allocation in that maximization. Differentiating (39) gives the first equality in (37). Removing winner i preserves all previously selected rival orders and fills exactly k_i^* additional ones. The regulated orders among them contribute between zero and k_i^* to exposure. Summing proves its bounds. The resulting uniform Lipschitz bound $\bar{q}n$ permits integration across breakpoints and interchange of expectation with integration, proving (38). $\square$

The identity explains why observed purchases are an incomplete measure of policy incidence. A losing regulated order can price an unaffected winner, causing receipts to fall without changing allocation. Conversely, a regulated winner can face a payment supported by unaffected demand or retention. The policy reaches the seller through the counterfactual demand that prices its sales.

One immediate consequence is bounded receipt transmission:

$$0 \leq W_n(q, 0) - W_n(q, \tau) \leq n\bar{q}\tau \quad (40)$$

For $n \geq 1$ and $\tau > 0$, define cumulative transmission relative to offered capacity by $\alpha_R = (n\tau)^{-1} \int_0^\tau \bar{\chi}_n(q, t) dt \in [0, 1]$. It is below one exactly when $\int_0^\tau [n - \bar{\chi}_n(q, t)] dt > 0$. Relative to actual sales, strict attenuation requires $\int_0^\tau \mathbb{E}[m - \chi_n] dt > 0$. Unregulated demand or retention must support a positive measure of payment slots in the corresponding comparison. The remaining question is how this loss varies with the stock the gate chooses to produce.

A worked payment comparison. Take $q = 2.5$, $n = 2$, $\bar{q} = 1$ and $H(y) = 6y$. Unaffected buyer A submits marginal values 16, 14; regulated buyer B submits $12 - \tau, 10 - \tau$; and unaffected rival C submits one value, 8. Buyer A receives both packages throughout, so actual retained stock is $y = 0.5$, worth 3, and regulated purchases are zero. Yet B initially sets both critical payments. Figure 3 shows how unaffected competition and retention successively buffer seller value.




Figure 3: Two buffers for VCG seller value at fixed inventory. Dots mark the with- C observations in Table 1.

Table 1: Critical orders, transfers and retention in the with- C case.

τ	Critical rival orders	Buyer payment	Actual retained value	Total seller value
0	Regulated 12 + regulated 10	22	3	25
3	Regulated 9 + unaffected 8	17	3	20
5	Unaffected 8 + regulated 7	15	3	18
7	Unaffected 8 + retention $\rho_1(0.5) = 6$	14	3	17

Retention in a critical order is a counterfactual opportunity cost, not an extra physical package kept in the actual allocation. The table therefore adds only $H(0.5) = 3$ to the payment. These buffers concern value at a *fixed inventory*: unchanged sales to A do not establish unchanged optimal extraction. That conclusion requires comparing the payment loss across inventories, as in the next subsection.

3.8 When the gate preserves production

Let $\Xi = (z, n)$ range over the full feasible implementation set $\mathcal{X}$. The regional cost is $C(z)$ and its output is $q(z)$. Profit is

$$\Pi(\Xi, \tau) = W_n(q(z), \tau) - C(z) \quad (41)$$

The selective downstream cost leaves regional primitives and the chosen equilibrium graph fixed. The gate can still change village coverage, prices, input margins, coercion, residual stock and offered capacity. Using the attained envelope $\mathcal{C}(q)$ collects these possibilities into a comparison across all attainable outputs and allowed capacities.

Let Ξ_0 be an attained baseline maximizer and define $\Delta_0(\Xi) = \Pi(\Xi_0, 0) - \Pi(\Xi, 0)$.

Theorem 3.11 (Global output resilience). *For every feasible competitor Ξ and policy τ ,*

$$\Pi(\Xi_0, \tau) - \Pi(\Xi, \tau) = \Delta_0(\Xi) - \bar{q} \int_0^\tau [\bar{\chi}(\Xi_0, t) - \bar{\chi}(\Xi, t)] dt \quad (42)$$

The baseline remains optimal exactly when the right-hand side is nonnegative for every feasible competitor. If it is strictly positive for every competitor with a different output, baseline physical output remains uniquely optimal. This conclusion permits multiple cost-minimizing nominal implementations of the same output.

Proof. Apply (38) separately to the two choices and subtract. Their implementation costs have no direct policy term. The resulting identity is (42); global optimality is precisely the collection of its comparisons. Restricting the strict comparisons to different-output competitors proves the quantity claim without imposing nominal uniqueness. $\square$

Production depends on relative profit. The original choice can remain optimal after a large receipt loss if alternatives lose as much or were sufficiently less profitable to begin with. It can cease to be optimal despite positive margins when another inventory is less exposed. The theorem combines market opportunities and the cost of Selva's feasible allocations in the comparison governing extraction.

Corollary 3.12 (Protection by unaffected payment-setting orders). *If $\bar{\chi}(\Xi_0, t) = 0$ for almost every $t \in [0, T]$, the baseline remains globally optimal throughout that interval. A unique baseline remains unique. If all baseline maximizers have the same output and at least one is protected, that output remains uniquely optimal.*

Proof. The protected baseline loses no receipts. Every competitor's receipts weakly decrease by (37), hence $\Pi(\Xi, \tau) \leq \Pi(\Xi, 0) \leq \Pi(\Xi_0, 0) = \Pi(\Xi_0, \tau)$. Any strict baseline comparison remains strict. Only baseline maximizers can remain tied with the protected choice, which proves the final statement. $\square$

Complete protection requires unaffected demand or retention to determine every critical payment almost surely over the policy interval. A high unaffected floor by itself allows regulated orders to set still higher payments and transmit the shock. Empirically, the relevant opportunities include the alternatives that set payments, as well as observed buyers and laundering volumes.

Corollary 3.13 (Falling rents with unchanged extraction). *Suppose Ξ_0 uniquely maximizes baseline profit, a set B contains Ξ_0 , and $\bar{\chi}(\Xi, t) = \kappa(t)$ for all $\Xi \in B$ and almost*

every $t \in [0, T]$. If (42) is strictly positive for all $\Xi \notin B$ throughout the interval, the baseline choice is unchanged and its profit falls by $\bar{q} \int_0^\tau \kappa(t) dt$. If there is a uniform outside gap $\delta = \inf_{\Xi \notin B} \Delta_0(\Xi) > 0$ and an economically bounded package count $\bar{Q}$, it suffices to restrict policy to

$$0 \leq \tau < \min\{T, \delta/(\bar{q}\bar{Q})\} \quad (43)$$

Proof. Inside B , the common exposure loss cancels from every profit difference. Outside, the assumed comparisons retain baseline superiority. The exposure bound gives an integrated differential loss of at most $\bar{q}\bar{Q}\tau$, so the strictly positive outside gap implies the stated sufficient interval. $\square$

With continuous output, nearby quantities can have arbitrarily close profits. Equal exposure preserves their ranking by imposing the same receipt loss on each. The global comparison handles more distant quantities, exit and alternative village networks. This is how selective regulation can lower the intermediary's rent while leaving its preferred production choice unchanged.

3.9 Partial transmission and changes in production

When exposure varies with inventory, the gate can adjust output in response to policy. Consider a fixed smooth feasible regime with regular interior optimum $z^*(\tau)$ and negative definite Hessian $\mathcal{H} = \Pi_{zz}$ in independent feasible coordinates. Let $a = \nabla_z q \neq 0$. On this branch, policy enters profit only through $W_n(q, \tau)$.

Theorem 3.14 (Local production transmission). *At such an optimum,*

$$\frac{dq^*}{d\tau} = -W_{q\tau} a^\top \mathcal{H}^{-1} a \quad (44)$$

Thus $W_{q\tau} < 0$ implies falling extraction. Where expected exposure is differentiable, define $\Lambda = -a^\top \mathcal{H}^{-1} a > 0$ and $\gamma = \bar{q} \partial_q \bar{\chi}_n$. Against a direct per-unit levy evaluated at the same zero-policy baseline,

$$(dq^*/d\tau)_0 = -\Lambda\gamma, \quad (dq^{\text{levy}}/dt)_0 = -\Lambda \quad (45)$$

Strictly partial local contraction follows when $0 < \gamma < 1$; zero local transmission follows when $\gamma = 0$.

Proof. Differentiate all the regional first-order conditions: $\mathcal{H}z_\tau + W_{q\tau}a = 0$. Multiply by $a^\top$ to obtain (44). Negative definiteness gives $a^\top \mathcal{H}^{-1}a < 0$. Differentiating the exposure identity gives $W_{q\tau} = -\bar{q} \partial_q \bar{\chi}_n = -\gamma$. A direct levy subtracts $tq(z)$ from the baseline profit, giving the same Hessian and a policy-gradient term $-a$ at zero. This yields the second derivative in (45) and the comparison. $\square$

It is the additional exposure of marginal stock that governs the local production response. The level bound $0 \leq \bar{\chi}_n \leq n$ places no restriction on that derivative. Buyer distributions

and retention technology can, however, determine its sign and magnitude, as the following example shows.

Corollary 3.15 (An analytical partial-transmission condition). *Consider one offered package and independent unaffected and regulated unit-demand values U, R . Let F_U be the unaffected distribution and f_R a continuous regulated density. On an inventory interval with differentiable reserve $\rho(q)$,*

$$\gamma = -\bar{q}\rho'(q)[1 - F_U(\rho(q))]f_R(\rho(q) + \bar{q}\tau) \quad (46)$$

Concave retention makes this weakly nonnegative. A strictly decreasing reserve, positive rival probability and density, and a product in (46) smaller than one give strictly partial local production transmission under the regular regional optimum conditions.

Proof. A regulated order occupies the critical slot exactly when $U > R - \bar{q}\tau > \rho(q)$. Hence expected exposure is $\int_{\rho(q) + \bar{q}\tau}^{\infty} [1 - F_U(r - \bar{q}\tau)]f_R(r)dr$. Differentiate its lower endpoint in q and multiply by $\bar{q}$ to obtain (46). Concavity makes $\rho' \leq 0$. The remaining conditions allow application of Theorem 3.14. $\square$

Several adjustments can consequently occur at different policy levels. Actual allocation changes who receives mineral; counterfactual allocation changes who sets its price; and a regional switch changes production, coverage or local implementation. Residual stock can also vary continuously while Q stays fixed. The appendix gives symbolic examples in which profitable sale survives a procurement switch and full-support uncertainty produces a smooth output response.


Selective regulation changes extraction when it changes the relative return to feasible inventories. Retention and unaffected competition can preserve that ranking, sometimes even while the seller's profit falls. They can also provide only partial protection. The distinction follows from the opportunities the intermediary loses at each inventory, rather than from whether regulated demand or laundering remains positive.

3.10 The institutional implication

The limitation of selective demand policy is therefore institutional. Some buyers remain outside its reach, and concealed origin permits mineral to reach additional outlets. Together these opportunities can continue to reward the same pattern of production in Selva.

Traceability addresses the connection between origin and market access. If obtaining a verified claim requires a verifiable production history, buyers can condition their demand on that history. The claim must have physical backing, and the participants controlling admission must prefer verification to capture. The next section derives those incentives. Its purpose is to change the opportunities through which illegal production can earn a return.

4 Provenance, private incentives, and the cost of credibility



Origin claims, scrutiny and collectible obligations change exterior market access.

The value of certification depends on events far removed from the buyer. A purchaser can pay for an admissible mineral without observing the mine, the conditions of extraction, or the actions of the inspector who approved it. Trade creates a return to a credible origin claim and, with it, a return to supplying the same claim on different material. The problem is to make the first return available while limiting the second.

Credibility requires both evidence and incentives to act on it. The resource cost of discovering privately observed facts is central to the costly-verification literature [Townsend, 1979]. The allocation of those costs and rewards also affects the reports that reach an authority. In a field experiment on environmental audits in Gujarat, Duflo et al. [2013] show that reforming auditor incentives improved reporting and reduced firms' pollution. Here the object of verification is the origin of a traded mineral. Its commercial life extends beyond the initial inspection, creating further opportunities to produce evidence and further interests in concealing it.

A persistent history connects those later decisions to the original transaction. If certification can be withdrawn after an adverse finding, the licensed producer has collateral and future business at stake; the inspector must be compensated for the consequences of a false report; and a downstream buyer may benefit from discovering a defect before its losses increase. The organizer must pay for the cooperation needed to admit the mineral and then decide whether to buy silence or bear the risk of discovery. Thus a continuing claim about origin changes the private cost of sustaining a false declaration.

Material can enter certified trade and later lose eligibility. Withdrawal changes its legal status but leaves its extraction history and the damage it caused intact. Verification also has costs even when it deters fraud: inspectors must be paid, evidence retained, disputes adjudicated and capital held against unresolved claims. We determine these

obligations from the physical claims being made and the incentives to investigate, report and cooperate in concealment.

4.1 Physical origin and persistent claims

A buyer acquires a claim about the mineral’s production history together with the mineral itself. Resale changes who bears the commercial exposure to that claim. It leaves the underlying history unchanged. Let $\mathcal{M}$ denote physically identified mineral portions and let m be their additive measure of standardized recoverable mineral content. A live claim c records a fresh identifier, owner, custody evidence, positive mass $v(c)$ and backing $\psi(c) \subseteq \mathcal{M}$, with $m(\psi(c)) = v(c)$. Different live claims have disjoint backing, up to zero mass. Inspection and custody maintain these physical-accounting conditions; cryptography authenticates records of the conditions.

Let S_t be the finite set of live claims after settlement event t . A mint admits inspected, physically registered material. A transfer consumes live inputs once and creates fresh outputs. Let m_{loss} be the recoverable mineral content lost in the transfer. The outputs satisfy

$$v(O) + m_{\text{loss}} = v(I), \quad \psi(O) \subseteq \psi(I), \quad m_{\text{loss}} \geq 0, \quad (47)$$

where $v(I) = \sum_{c \in I} v(c)$ and $\psi(I) = \bigcup_{c \in I} \psi(c)$; outputs are disjoint and recorded losses cannot back another claim. Redemption consumes a claim without creating a successor. Redemption may release material for use or record destruction, so retiring a claim does not always destroy mineral. Inputs, outputs and registration locks update atomically. Signatures authenticate authorized actions and sound proofs verify the specified record relations. The required physical facts remain separate premises.

Each mint has a persistent physical-origin identifier. New claims consume only previously live inputs, making the accepted transaction history a directed acyclic graph. Write $\text{Anc}(c)$ for the origin mints recorded for claim c . Under the conservative lineage rule, a mint records its own origin and each transfer output inherits the union of its inputs’ origins. Splitting, merging, and refining preserve a path back to admission. Supply-chain and provenance standards provide representations for event records and derivation links [GS1, 2023, Moreau and Missier, 2013]. Establishing the truth of the physical declaration remains an inspection task. Where records are confidential, an authorized opening procedure gives the relevant investigator access to evidence; public proof verification can preserve that confidentiality.

The source of illegal stock determines where illegal-mining leakage can begin. Index countries by *true extraction origin* k , and fix the production cohort. Let $\mathcal{I}$ contain the countries that supply illegally mined material in this cohort. Available illegal stock $q_{D,k}$ and attempted diversion d_k satisfy

$$0 \leq d_k \leq q_{D,k}, \quad k \notin \mathcal{I} \implies q_{D,k} = d_k = 0. \quad (48)$$

The restriction concerns illegal extraction. False measurement and missing custody evidence create other eligibility failures. An exporting location or an absence of reported cases does not establish $q_{D,k} = 0$: the stock restriction is about true origin.

Trade can carry material from an illegal mine to a foreign refinery, certification facility, or purchaser. Its true origin remains k at each stage. Limiting fraudulent admission to the mining country would require an origin-country admission rule together with controls against substitution and foreign reminting. Historical material likewise retains its cohort's production history after illegal mining ceases at that origin. Purchasers can therefore face illegal-origin exposure in countries with no domestic illegal extraction.

4.2 Certification withdrawal without deleting history

Maintain the transaction graph separately from current certification status. Let $\mathcal{O}_t^{\text{rev}}$ be the origins whose certification has been revoked by an effective final adverse decision, and $\mathcal{O}_t^{\text{susp}}$ those under effective temporary suspension. The certified live set is

$$C_t = \{c \in S_t : \text{Anc}(c) \cap (\mathcal{O}_t^{\text{rev}} \cup \mathcal{O}_t^{\text{susp}}) = \emptyset\}. \quad (49)$$

An authenticated authority records each decision, its scope, effective settlement event and evidence reference. Decisions and settlements share one ordered sequence: every certified settlement validates against the status state immediately preceding its commitment. A pending challenge can therefore block settlement while evidence is examined. A final decision withdraws certification without requiring the current owner's consent. An appeal changes status through a new authenticated decision under a specified reinstatement rule; it never erases the history.

The institution withdraws access to certified settlement and its associated contractual privileges. Physical custody, notice, segregation, lawful unverified disposition, and civil recovery follow their own procedures; certification authority does not confer a universal power to confiscate property. Material already redeemed or consumed has left S_t . Its history can support contractual or recall claims, although there is no longer a live certified claim to remove.

Exclusion follows the physical origin and its derived material through redemption, relabelling, and attempted readmission. Intake checks prevent excluded backing from returning with a new identifier, except after an authorized reinstatement. Reliable physical identification and access to the history make this rule effective. A concealed substitution that defeats these controls is a new enforcement failure. A custody defect arising later in the chain instead attaches to that event and its affected descendants, leaving an innocent origin's unrelated branches eligible.

Proposition 4.1 (Conservation and retrospective exclusion). *Suppose the initial claims have disjoint, correctly measured backing and finite genesis histories. Maintain physical identification, complete recorded ancestry, the atomic transition rules above, authenticated*

status decisions and intake checks against excluded physical origins. Then every live claim has finite provenance, no positive mass backs two live claims, and

$$v(S_t) = m\left(\bigcup_{c \in S_t} \psi(c)\right) = v(S_0) + A_t - X_t, \quad (50)$$

where A_t is cumulative admitted claim mass and X_t is cumulative recorded processing loss and redemption mass. A status decision alone changes neither side of this physical-backing identity. After an origin's revocation becomes effective, and until any authorized reinstatement, none of its live recorded descendants can settle as certified or regain certification through a fresh identifier.

Proof. Induct on the ordered events. A mint adds measured backing disjoint from all surviving claims. A transfer replaces its inputs with disjoint portions of that backing and records the residual as loss. Redemption removes backing from the live claim account. These operations preserve unique backing and give the mass identity by additivity and telescoping. If previously redeemed material is validly readmitted, its mass enters A_t again, matching its earlier exit in X_t ; the identity is a live-stock account, not a count of distinct historical extraction. Fresh identifiers and previously live inputs preserve an acyclic finite history. Status events alter eligibility only.

At an effective revocation, every existing claim containing that origin in its ancestry is excluded by (49). Every successor inherits that ancestry. Induction therefore excludes all later recorded descendants at the status check governing their settlement. The ordered validation rule prevents a concurrently submitted transfer from using a stale status. Intake checks prevent the same excluded physical origin from escaping through another mint. No subsequent event changes this conclusion without the specified reinstatement decision. $\square$

The proposition explains what persistent history accomplishes once a defect is effectively discovered: transfers preserve the reach of exclusion. Initial truthfulness and the probability of discovery are determined by the incentive problems below. Live mass affected by final origin revocations is counted once:

$$Q_R(t) = \sum_{c \in S_t} v(c) \mathbf{1}\{\text{Anc}(c) \cap \mathcal{O}_t^{\text{rev}} \neq \emptyset\}. \quad (51)$$

A claim with several revoked ancestors contributes only once. Summing historical transaction masses would repeatedly count the same mineral.

The precision of physical attribution determines the reach of a withdrawal. Exact attribution permits the affected portions to be separated from verified admissible backing. Where separation is unavailable, the conservative rule in (49) excludes the associated whole lot. Mixing can then expose admissible material to the loss of certification as well. The commercial loss exceeds the original illegal mass without increasing physical contamination: a recorded transformation link alone does not identify individual molecules.

For geographic aggregation, let V_k be live certified mass physically attributable to true origin k , and U_k its illegally mined component. These analytical quantities partition physical mass even when operational records support only conservative lot exclusion. Hence $U_k = 0$ outside $\mathcal{I}$. Define ι_k^{live} as the illegal share of expected live certified mass from origin k and $\iota_{\text{world}}^{\text{live}}$ as the corresponding world share. For positive expected total volume,

$$\begin{aligned}\iota_{\text{world}}^{\text{live}} &= \frac{\sum_k \mathbb{E}[U_k]}{\sum_k \mathbb{E}[V_k]} = \sum_{k \in \mathcal{I}} w_k \iota_k^{\text{live}} \leq \sum_{k \in \mathcal{I}} w_k, \\ w_k &= \frac{\mathbb{E}[V_k]}{\sum_h \mathbb{E}[V_h]}, \quad \iota_k^{\text{live}} = \frac{\mathbb{E}[U_k]}{\mathbb{E}[V_k]}.\end{aligned}\tag{52}$$

Each origin enters the world rate through its share of expected live mass. Origins with zero expected volume are omitted because their rates are undefined. For exposed origins, diversion and enforcement determine surviving leakage, which can be zero. This accounting makes the geographic incidence of exposure depend on the actual composition of supply.

4.3 The horizon of collectible exposure

The commercial life of an origin claim can outlast its first owner and inspection committee. Deterrence requires a discovery and claims horizon T_H , an assignment of responsibility for older vintages, and assets from which adverse awards can be collected. Retained collateral, continuing licence value, or funded runoff insurance can support those awards under their contractual terms. Evidence discovered after financial support expires can still justify decertification, but may yield little recovery from an earlier participant. The duration of the record and the duration of financial exposure are separate institutional choices.

Several vintages may draw on participant i 's bond B_i . The resulting simultaneous collectible payments cannot exceed B_i , so the institution must assess joint deviations and claims priority against the common pool. It can reuse or release the bond after the corresponding exposure closes or replacement funding is secured. Retaining evidence, delaying collateral release, and maintaining insurer capital consume resources throughout this period. These continuing costs make retrospective exclusion credible before any adverse finding occurs.

4.4 Why a purchaser investigates and reports

A downstream purchaser may have a reason to investigate a claim approved by the organizer, licensed custodian, and original inspectors. Withdrawal can interrupt its operations or leave it with a contractual loss. Acting earlier may reduce that loss. The purchaser can investigate directly or pay an independent inspector, insurer, or specialist. In either case the activity must produce evidence that the authority can use to change eligibility. Repeated consultation of the same unverified certificate supplies no new physical observation.

Consider a purchaser j with posterior probability π_j that an exposed claim has an illegal origin. The purchaser chooses effort $e_j \geq 0$ at cost $c_j(e_j)$, with $c_j(0) = 0$. Conditional on fraud, the probability of obtaining actionable evidence early is $s_j(e_j)$, with $s_j(0) = 0$. Evidence is conclusive in this benchmark and has no false positives. Its production is independent of later evidence from other channels, conditional on fraud and the public state. Let P_{-j} denote the probability of discovery through those other channels over the remaining horizon.

Let $L_{E,j}$ be the purchaser's net private loss after its own early finding and report, $L_{L,j}$ its loss after later discovery elsewhere, and $L_{U,j}$ its loss when the defect remains undiscovered during the horizon. All three are currency losses relative to an otherwise identical eligible purchase. They include lost certification services, operational consequences, and contractual payments or recoveries actually available to the holder. Timing and discounting are incorporated in these amounts. We hold purchased quantity and price fixed; the value of refusing acquisition would enter a pre-purchase screening decision separately.

Define the advantage of early reporting by

$$A_j = P_{-j}L_{L,j} + (1 - P_{-j})L_{U,j} - L_{E,j}. \quad (53)$$

Concealment leaves the specified later-discovery law and losses unchanged. This fixes the purchaser's alternative to disclosure. Where concealment changes liability, evidence access, or resale opportunities, those changes must enter that alternative payoff.

Proposition 4.2 (Investigation and voluntary disclosure). *After conclusive evidence of fraud, reporting is optimal exactly when $A_j \geq 0$, with indifference at equality. When reporting is chosen, the gain from investigation over zero effort is*

$$\pi_j s_j(e_j) A_j - c_j(e_j). \quad (54)$$

Suppose $A_j > 0$, the feasible effort interval is compact, both primitives are continuous, s_j is increasing and concave, and c_j is increasing and strictly convex. The optimal effort is unique. If the primitives are differentiable, an interior optimum satisfies

$$c'_j(e_j) = \pi_j A_j s'_j(e_j). \quad (55)$$

With twice differentiable primitives, a positive marginal success probability, and $c''_j - \pi_j A_j s''_j > 0$ at the optimum, its response to A_j , holding beliefs and other channels fixed, is

$$\frac{de_j}{dA_j} = \frac{\pi_j s'_j(e_j)}{c''_j(e_j) - \pi_j A_j s''_j(e_j)} > 0 \quad (56)$$

when $\pi_j > 0$. At boundaries the optimal effort is weakly increasing in A_j . A fixed startup cost f_j induces entry only if the maximized variable gain in (54) covers f_j .

Proof. Given conclusive evidence, reporting costs $L_{E,j}$; concealment has expected cost $P_{-j}L_{L,j} + (1 - P_{-j})L_{U,j}$. Subtraction gives (53). Before investigation, these alternatives

differ only when the claim is defective and early evidence is produced. Multiplying by that probability and subtracting effort cost yields (54). Its strict concavity establishes uniqueness; its first derivative gives (55). Implicit differentiation gives (56). Increasing differences in (e_j, A_j) give weak monotonicity including corners: adding the two optimality inequalities for a putative lower effort at a larger A_j contradicts increasing s_j , unless the choices coincide. Finally, an entrant compares its maximized gain less f_j with zero. $\square$

The return to diligence comes from a loss that information allows the purchaser to avoid. An early finding may prevent processing and resale of a defective batch or allow recovery before a liable supplier's assets dissipate. Enforceable liability for late disclosure can create the same incentive. Exposure alone is insufficient. If $L_{E,j} = L_{L,j} = L > 0$ and $L_{U,j} = 0$, then $A_j = -(1 - P_{-j})L$. For $P_{-j} < 1$, discovery merely accelerates a loss that might never have occurred. A holder with no other private benefit from evidence chooses no costly investigation.

Insurance and recourse shape this comparison. An indemnity can make early disclosure attractive; complete coverage of avoidable losses can also remove the insured party's incentive to investigate. Deductibles, diligence conditions, or investigation undertaken by the insurer can restore a return to evidence. A reporting reward requires both funding and a contract under which manufacturing fraud to claim the reward is unprofitable. The relevant economic choices concern the allocation of losses and information duties.

Voluntary diligence also responds to the success of deterrence. At zero posterior fraud risk, (54) gives zero effort when investigation is costly and has no separate benefit. Continuing scrutiny in a faithful equilibrium therefore requires committed paid inspection, an independently valuable diligence service, or an explicit error or off-path environment supporting investigation. The market solved below instead sustains positive private investigation through positive on-path risk.

4.5 Independent evidence and the organization of scrutiny

For a fixed fraudulent origin, let r_0 be the probability of effective baseline forensic discovery by the end of a specified horizon. Let r_j be the conditional probability that downstream channel j produces evidence, reports it and obtains an effective adverse decision. Thus r_j includes the reporting incentive and evidentiary requirements; it is not just the probability that somebody looks at a record. Conditional independence yields

$$P_n = 1 - (1 - r_0) \prod_{j=1}^n (1 - r_j). \quad (57)$$

The first effective discovery removes the affected claim. Later findings may strengthen the case or add evidentiary costs. Financial and physical accounting still count the same bond and the same withdrawn quantity once.

Proposition 4.3 (Conditional survival of an admitted fraud). *Under the conditional independence and effective-decision requirements, adding a channel changes discovery prob-*

ability by $(1 - P_n)r_{n+1}$. If $r_j \geq \underline{r} > 0$, the probability that an admitted fraudulent origin remains undiscovered is at most $(1 - r_0)(1 - \underline{r})^n$.

Proof. The probability of no effective discovery is the product of the failure probabilities. Subtract the expressions for n and $n + 1$ channels to obtain the increment. Bounding each factor by $1 - \underline{r}$ gives the result. $\square$

Independent evidence reduces the probability that an admitted fraud survives. The source of that improvement is investigative effort with access to distinct evidence. To see the difference between this effort and the number of reviewers, suppose $r_j = 1 - \exp(-\alpha e_j)$. Equation (57) becomes

$$P_n = 1 - (1 - r_0) \exp \left(-\alpha \sum_j e_j \right). \quad (58)$$

Only total effort enters this expression. Dividing the same effort among more reviewers leaves discovery unchanged. With payoff $-c_j(e_j) - \pi L \prod_k [1 - s_k(e_k)]$, a reviewer shares the benefit of other parties' investigations, which reduce its own marginal return to effort. Information sharing can consequently economize on investigation and induce free riding at the same time.

The organization of evidence supplies a second limit. Several institutions may depend on the same provider or adjudicator. If they share a blind state with probability z , and discovery outside that state obeys (57), unconditional survival is $z + (1 - z)(1 - r_0) \prod_j (1 - r_j)$. Additional independent channels outside the blind state leave its probability unchanged. Independence must therefore refer to effective evidence and enforcement; additional ledger nodes alone do not supply it.

4.6 Finite diversion opportunities and distributed discovery

An organizer compares the return to false certification with the value of selling its mineral through the available dirty-market outlet. It must also secure the cooperation of agents whose reports control admission. This is a problem of joint as well as individual incentives: side agreements can alter the information transmitted through an organization [Tirole, 1986]. Here those agreements must be financed before fraudulent receipts arrive, and later holders can acquire evidence outside the original approval group.

Fix a market state Ω , a committed dirty-stock continuation $W_D(\cdot; \Omega)$, and an organizer holding q units. A feasible action a specifies all fraudulent batches, custody paths, and false approvals in the exposure pool, with total diversion $0 < d(a) \leq q$. Let $F_C(a; \Omega)$ be the joint certified receipts net of ordinary payments maintained for cooperating agents, and let $k_f(a; \Omega)$ be additional physical and concealment costs, excluding the side payments and sanctions below. The gain before capture expenses is

$$G_\Omega(a) = F_C(a; \Omega) - k_f(a; \Omega) - [W_D(q; \Omega) - W_D(q - d(a); \Omega)]. \quad (59)$$

Diversion yields $F_C - k_f + W_D(q - d(a); \Omega)$ instead of $W_D(q; \Omega)$. The acquisition bill for the original stock is already paid and cancels. What remains is a finite opportunity cost: the loss in the continuation value of the dirty stock. This value is optimized jointly over the remaining stock, so separately optimized batch opportunities generally cannot be added. The state Ω retains the original stock, announced menus, clearing rule, and observed history. For a mechanism Σ , therefore, $W_D(y; \Omega) = W_\Sigma(y \mid q, \mathcal{H})$. Inventory-contingent reoptimization is available only where the contingency itself was committed and verified. The comparison applies to large diversions as well as marginal ones and requires neither differentiability nor concavity.

The original approval group comprises the organizer, licensed producers or custodians, and assigned inspectors. Let S contain the organizer, indexed by 0, and each original participant whose approval action a requires. A finite, disjoint set J contains downstream buyers, insurers, or investigators able to obtain and report actionable evidence. These parties need not sign the mint. The organizer can leave a downstream channel active and accept discovery risk; its entry adds a source of evidence without necessarily adding a veto over admission.

At the offer stage, the organizer knows which channels in J it can approach and their reservation payoffs. A silence agreement is enforceable under the specified side-contract protocol; inaccessible later reporters remain part of residual discovery. Private types, unobservable investigations, or sequential bargaining would add their own information and incentive constraints to feasible offers. The current comparison fixes this access and commitment structure. Effective discovery in turn requires authorized evidence access, disclosure, adjudication, and withdrawal of certification.

4.7 The funded cost of admission and suppression

The price of cooperation depends on losses that can actually be imposed. For a silence set $T \subseteq J$, let $a_i(T) \geq 0$ be original participant i 's expected incremental collectible loss relative to honesty under the remaining discovery channels. This includes exposed collateral, lost continuation rents, and enforceable net recourse over the claims horizon. A common bond is counted once across unresolved batches and vintages; current losses already deducted in (59) are excluded. A buyer's inventory loss affects the organizer's liability only through an actual contractual or institutional obligation.

For $j \in T$, let $\tau_j(T) \geq 0$ be the least delivered payment that makes silence acceptable relative to the reporter's actual disclosure alternative, holding the other offers fixed. The side-contract protocol makes these reservation payments implementable, including the consequences of deviating to disclosure. All parties are risk neutral, retain ordinary contractual earnings, have no omitted benefit from fraud, and accept at equality. A gross transfer t_i delivers $\lambda_i t_i$, where $0 < \lambda_i \leq 1$. Losses and reservation payoffs depend on the silence set but remain unchanged by transfers above the minimum. Offers, false approvals, and purchased silence are jointly enforceable over the modeled horizon. This commitment

prevents a paid investigator from reporting anyway or reopening the bargain later; either option would change the acceptance constraints.

Let $F_0(a, \Omega)$ be cash or enforceable credit available before settlement, after other committed advances. Future fraudulent receipts count only if they can actually be pledged. Define

$$\mathcal{B}(T) = \sum_{i \in S \setminus \{0\}} \frac{a_i(T)}{\lambda_i} + \sum_{j \in T} \frac{\tau_j(T)}{\lambda_j}, \quad (60)$$

$$K_J^F(a) = \min_{\substack{T \subseteq J \\ \mathcal{B}(T) \leq F_0(a, \Omega)}} \{a_0(T) + \mathcal{B}(T)\}, \quad (61)$$

The value is $+\infty$ when no plan can be financed. Expected future liability $a_0(T)$ reduces the return to the action but supplies no cash for its advances. When a person performs several roles, its joint acceptance constraint replaces the separate entries and its assets are counted once.

Theorem 4.4 (Funded capture with optional suppression). *Under this protocol, the action a has a strictly profitable funded implementation if and only if $G_\Omega(a) > K_J^F(a)$. Consequently no profitable funded diversion exists when $G_\Omega(a) \leq K_J^F(a)$ for every feasible nonzero joint action. Strict inequalities $G_\Omega(a) < K_J^F(a)$ exclude indifferent implementations as well.*

Proof. For a fixed T , original recipient i gains $\lambda_i t_i - a_i(T)$ relative to honesty. Acceptance therefore requires $t_i \geq a_i(T)/\lambda_i$. Each purchased reporter analogously requires $t_j \geq \tau_j(T)/\lambda_j$. The joint implementability assumption makes these bounds sufficient as well as necessary. Thus every accepted schedule implementing T costs at least $\mathcal{B}(T)$ before settlement. If that bill exceeds F_0 , no such schedule is funded. Otherwise equality payments are funded and leave organizer surplus $G_\Omega(a) - a_0(T) - \mathcal{B}(T)$; larger transfers cannot improve it. Since J is finite, the minimum over nonempty funded choices is attained. Maximizing organizer surplus is therefore exactly subtracting $K_J^F(a)$ from $G_\Omega(a)$. The empty case has no funded implementation. Applying the same comparison to every joint action proves both deterrence statements. At equality the organizer and recipients can be indifferent, so exclusion then requires an explicit honest tie rule. $\square$

Proposition 4.5 (A funded monotonicity condition). *Add a channel $j_{\text{new}} \notin J$, holding the action, gain and funding cap fixed. Suppose every expanded funded plan T' projects to an admissible old plan $T = T' \cap J$ satisfying*

$$\mathcal{B}_J(T) \leq \mathcal{B}_{J \cup \{j_{\text{new}}\}}(T'), \quad a_{0,J}(T) + \mathcal{B}_J(T) \leq a_{0,J \cup \{j_{\text{new}}\}}(T') + \mathcal{B}_{J \cup \{j_{\text{new}}\}}(T').$$

Then $K_{J \cup \{j_{\text{new}}\}}^F(a) \geq K_J^F(a)$.

Proof. The first inequality and unchanged cap make each projection funded in the old game. Its total cost is at least the old minimum; the second inequality makes the expanded plan at least as costly. Minimizing over expanded plans proves the result, including an

empty expanded feasible set. An empty old feasible set precludes any expanded funded plan by projection. $\square$

The projection condition has a direct economic interpretation. Adding an investigator weakly increases the cost of capture if it leaves old transfer efficiencies unchanged, weakly raises the original participants' losses and existing reporters' reservation payments under each projected plan, and adds a nonnegative silence bill. Independent discovery with fixed conditional liabilities supplies the original-loss comparison. Reporting reservations, funding, and feasible agreements must also satisfy the stated conditions when they respond to entry.

4.8 Entry and the limits of suppression

Consider the transparent subclass with independent effective discovery probabilities r_0 for baseline enforcement and r_j for channel j . All original participants face the same discovery event with fixed net collectible conditional losses $D_i \geq 0$. Put

$$D_G = D_0 + \sum_{i \in S \setminus \{0\}} D_i / \lambda_i, \quad P(T) = 1 - (1 - r_0) \prod_{j \notin T} (1 - r_j).$$

Take fixed gross silence costs $c_j^{\text{sil}} \geq 0$, derived from reporting reservations, and assume the advance funding constraint does not bind. The minimum cost becomes

$$K_n = \min_{T \subseteq \{1, \dots, n\}} \left\{ \sum_{j \in T} c_j^{\text{sil}} + D_G P(T) \right\}. \quad (62)$$

In this subclass, adding an independent channel weakly increases K_n . Deleting a newly suppressed channel saves its nonnegative bill; deleting an active one weakly reduces discovery. Either deletion produces an old feasible candidate at weakly lower cost. The same reasoning gives funded monotonicity when the original advance bill is $P(T) \sum_{i \neq 0} D_i / \lambda_i + \sum_{j \in T} c_j^{\text{sil}}$ and the funding cap is unchanged. Holding opportunities and the specified reporting behavior fixed, the set of profitable frauds therefore contracts weakly.

The comparison depends on collection. Fixed D_i requires retained security, solvent recourse, or another specified payment source throughout the claims horizon. Historical traceability can identify responsibility after a licensee's assets have been exhausted, but identification then supports little financial recovery. Transfers among cooperating parties also enter their joint payoff once. These restrictions give economic content to the conditional losses multiplied by discovery probabilities.

Proposition 4.6 (Homogeneous channels and saturation). *If $r_j = r \in (0, 1)$ and all channels have the same gross silence cost $c_j^{\text{sil}} = c^{\text{sil}} \geq 0$, then*

$$K_n = \min\{D_G P_n, n c^{\text{sil}} + D_G r_0\}, \quad P_n = 1 - (1 - r_0)(1 - r)^n. \quad (63)$$

In particular $K_n \leq D_G$. If $c^{\text{sil}} > 0$, $K_n \rightarrow D_G$; if $c^{\text{sil}} = 0$, $K_n = D_G r_0$ for every n .

Proof. Suppressing m channels costs $f(m) = mc^{\text{sil}} + D_G[1 - (1 - r_0)(1 - r)^{n-m}]$. Its continuous extension satisfies

$$f''(m) = -D_G(1 - r_0)[\log(1 - r)]^2(1 - r)^{n-m} \leq 0.$$

Concavity makes every value on $[0, n]$ at least the smaller endpoint, so the integer minimum is attained at $m = 0$ or $m = n$, proving the formula. The first endpoint is at most D_G and converges to it. For $c^{\text{sil}} > 0$ the second eventually exceeds D_G , whereas for $c^{\text{sil}} = 0$ it is the smaller endpoint. The case $n = 0$ follows directly. $\square$

The formula exposes a limit to deterrence through additional channels: the organizer can accept discovery and bear the collectible loss. With a binding advance cap, one endpoint may be unavailable, so the general comparison in (61) governs the choice. The closed form above applies to the nonbinding subclass.

The cost of a stipulated concealment target behaves differently. To keep discovery at most δ , where $r_0 \leq \delta < 1$, the number u of unsuppressed homogeneous channels must satisfy $(1 - r)^u \geq (1 - \delta)/(1 - r_0)$. Taking logarithms and dividing by the negative number $\log(1 - r)$ gives

$$u \leq u_{\max} := \left\lfloor \frac{\log[(1 - \delta)/(1 - r_0)]}{\log(1 - r)} \right\rfloor. \quad (64)$$

If each suppressed channel costs at least $c_{\min}^{\text{sil}} > 0$, maintaining this target costs at least $c_{\min}^{\text{sil}} \max\{0, n - u_{\max}\}$. A target below r_0 requires compromising baseline enforcement. Keeping fraud concealed can thus require an increasing suppression bill even though the organizer can choose instead to bear bounded expected sanctions. Distinguishing those two objectives explains the saturation result.

A common institutional failure with probability z leaves undetected survival at $z + (1 - z)(1 - r_0) \prod_{j \notin T} (1 - r_j)$. Independent channels outside that event cannot remove the common floor. Common ownership or shared evidence access can similarly make several nominal auditors a single effective channel. The number relevant for deterrence is therefore the number of distinct discovery opportunities under the stated protocol.

Deterrence is consistent with cooperative stability. In a transferable-utility specialization with honesty available, let a coalition's value be the larger of zero and its feasible fraud surplus. If every such surplus is nonpositive, all coalition values are zero and the zero incremental allocation lies in the core: it is efficient and no coalition can obtain a strict improvement. The capture condition identifies profitable funded departures from honesty directly.

4.9 Admission, retrospective removal, and surviving leakage

A certification institution affects both the amount of illegal material admitted and the duration of its eligibility. Let A be physical mass admitted in a cohort and $U \leq A$ its illegally mined component. Let $R_U \leq U$ and $R_A \leq A - U$ denote illegal and admissible

mass, respectively, that later lose certified status. Initially hold all other entries and exits fixed at zero. Contamination at admission and in the remaining certified stock is measured by

$$\iota^{\text{in}} = \frac{\mathbb{E}U}{\mathbb{E}A}, \quad \iota^{\text{live}} = \frac{\mathbb{E}(U - R_U)}{\mathbb{E}(A - R_U - R_A)}, \quad (65)$$

whenever the respective denominators are positive. These ratios aggregate expected physical masses. They are event probabilities only with the corresponding equal-mass, all-or-nothing convention, and generally differ from expected within-batch ratios. Total withdrawal $\mathbb{E}(R_U + R_A)$ and holders' losses describe further consequences of the same decisions.

Proposition 4.7 (The removal effect). *For a fixed positive-mass cohort with initial illegal share $\iota \in [0, 1]$, suppose only illegal units are withdrawn, each with probability P . There are no other exits. If $1 - \iota P > 0$, then*

$$\iota^{\text{live}} = \frac{\iota(1 - P)}{1 - \iota P} \leq \iota, \quad \iota - \iota^{\text{live}} = \frac{\iota P(1 - \iota)}{1 - \iota P}. \quad (66)$$

The inequality is strict for $0 < \iota < 1$ and $P > 0$. Holding these conditions fixed, live leakage weakly increases with admitted contamination and decreases with discovery:

$$\frac{\partial \iota^{\text{live}}}{\partial \iota} = \frac{1 - P}{(1 - \iota P)^2}, \quad \frac{\partial \iota^{\text{live}}}{\partial P} = -\frac{\iota(1 - \iota)}{(1 - \iota P)^2}. \quad (67)$$

Proof. Normalize admitted mass to one. Expected surviving illegal mass is $\iota(1 - P)$ and surviving total certified mass is $1 - \iota P$. Division, subtraction and differentiation give the displayed formulas. The independence of removal events across units is unnecessary for these expectations; the common marginal removal probability suffices. $\square$

The endpoints explain why continued trade matters for interpreting this rate. If all admitted units are illegal, partial withdrawal leaves a fully contaminated surviving stock. Full withdrawal leaves no stock on which to define a conditional rate. Conversely, zero illegal admission gives zero illegal-origin leakage whenever positive trade continues.

Whole-lot withdrawal changes the comparison because it may also remove admissible material. Consider a cohort of mass 100 containing 10 illegal units. Withdraw a mixed lot with 9 illegal and 89 admissible units, leaving an undetected illegal unit from a different, still-unrevoked origin. The remaining stock contains 1 illegal unit among 2 units. Illegal mass falls while its share rises from 10% to 50%. Proposition 4.7 excludes this admissible-mass removal. The example shows why the scale of remaining physical exposure and the precision of withdrawal belong beside the contamination rate in an assessment of certification.

In an ongoing market, (65) also accounts for redemption, consumption, physical loss, new cohorts, and reinstatement. A finding can create losses on goods already used, even though those goods have left the live stock. A bound on the probability of some ancestral failure

describes initial contamination risk. Strategic admission and subsequent withdrawal are the choices that determine how much illegal material remains certified.

4.10 Who pays for credibility

Credibility requires an organization that can collect evidence and act on it. In mineral supply chains, due diligence entails record maintenance, monitoring, supplier adjustments, and audits, with costs distributed across participants. Shared schemes and recognition of existing audits can reduce duplication [OECD, 2021]. The economic question is how this continuing service is financed and whether the parties whose cooperation it needs prefer to participate. The following accounts separate those incidence questions without treating the cited evidence as a calibration.

For a period with successful admitted volume V , unit fee ϕ and committed outside funds E_0 , an operator budget is

$$\begin{aligned} \phi V + E_0 \geq & C_{\text{admission}} + C_{\text{custody}} + C_{\text{records}} + C_{\text{forensic}} + C_{\text{adjudication}} \\ & + C_{\text{availability}} + C_{\text{capital}} + C_{\text{promised obligations}}. \end{aligned} \quad (68)$$

Each entry covers a distinct operator expense over the stated horizon, including compensation for honest effort and participation. Capital principal remains an asset, while the cost of retaining it is an expense. Promised obligations include rewards, compensation, insurance premia, or reserve funding under a consistent cash-flow convention. If the operator pays an insurer to cover a claim, the covered payment is the insurer's expense and is not charged again to the operator. Contingent solvency additionally requires collectible capital or committed finance adequate to the specified obligations.

Buyers pay for diligence and for the costs of holding inventory, replacing goods, interruptions and administering claims left with them by the contract. Licensees may bear deductibles, legal defense costs, or continuing bond charges. The operator's required fee finances its own expenses; costs borne elsewhere enter those parties' participation decisions. Each must receive a net payoff, after its costs and residual liabilities, at least as large as its outside option. Deterring false admission therefore addresses only one part of sustaining an operating market.

Proposition 4.8 (Cost of prevention and a scale limit). *Suppose an operating institution requires positive recurring expenditure $F > 0$ even when no fraudulent claim is found. It cannot finance that expenditure exclusively through seizures or discovery rewards that have zero expected receipts under perfect fidelity. Moreover, suppose total certified expenditure is at most $I < \infty$, fee revenue is φE_C with $0 \leq \varphi < 1$, outside funding is at most $\bar{E}_0$, and physical verification costs at least cV with $c > 0$. Any funded operating volume satisfies*

$$V \leq \frac{\varphi I + \bar{E}_0 - F}{c} \quad (69)$$

when $F + cV$ is a lower bound on operator cost. A negative numerator makes positive operation infeasible under these bounds.

Proof. Under perfect fidelity the proposed contingent revenue has expectation zero, while $F > 0$ is due; absent other funds the budget fails. For the scale bound, solvency and the revenue ceiling imply $F + cV \leq \varphi E_C + E_0 \leq \varphi I + \bar{E}_0$. Rearrange. Additional nonnegative costs can only tighten the bound. $\square$

The scale bound links the growth of certified trade to verification productivity and its funding base. Greater productivity, growing expenditure, or additional committed finance can relax it. With bounded fee receipts, continued growth in physical volume requires a way to service both current verification and historical exposure at declining resource cost. Fewer successful frauds alone provide no such productivity gain.

The distinction between private exposure and resource use is equally important for welfare. Damages, indemnities, seized bonds, and fees transfer resources among parties and alter incentives and liquidity. Investigator and lawyer time, interrupted services, replacement processing, and the opportunity cost of capital use real resources. Counting a lost premium, a compensating damages transfer, and a replacement invoice as three independent social losses would count parts of the same transaction again.

An institution can spend more on prevention while incurring fewer incidents and less litigation. Whether total resource cost rises depends on the relative changes in these components and on who bears them. Continuing capacity and private diligence remain part of the comparison after credibility improves. The next market makes that comparison explicit by solving fraud and investigation jointly under funded contracts.

4.11 A funded market with endogenous scrutiny and fraud

Fraud makes investigation valuable, and investigation reduces the return to fraud. A market must reconcile both choices with the beliefs on which purchasers act. We solve this interaction in a bounded posted-contract market with one unit of committed service demand. Procurement invoices are fixed, and the participation conditions determine whether all parties are willing and able to trade at those invoices.

4.11.1 Contracts, information and the physical opportunity

There is a unit continuum of equal-sized certification slots from one true extraction country in $\mathcal{I}$. Each slot can receive an admissible unit at certified quote h . Clean suppliers can serve this bounded demand at that price, which covers their costs and reservation returns. An organizer also controls one previously acquired illegal unit per slot. Its committed alternative market pays p_D per unit; thus the dirty continuation is $W_D(q) = p_D q$ on this example's operating range. Replacing clean procurement with that illegal unit sacrifices p_D , saves h , and incurs forgery cost c , whose cross-sectional distribution is uniform on

$[0, C]$. A buyer is randomly assigned to a slot independently of this privately observed cost. The gain before capture costs is

$$G(c) = \Gamma - c, \quad \Gamma = h - p_D. \quad (70)$$

The substitution saves clean procurement expenditure and sacrifices the committed dirty sale. Sunk acquisition payments cancel. Ordinary licensed-agent compensation is maintained, and clean production displaced before procurement has no unpaid contractual claim. A displacement charge, if one were owed, would reduce Γ .

The following amounts are per unit of slot mass. Each slot has a separate organizer and set of required origin participants, with no shared collateral across slots. Its organizer posts bond B_0 . The required licensed custodian and assigned inspectors post distinct bonds b_i , maintained until the entire cohort's claims horizon closes. Their sum with B_0 is $D = B_0 + \sum_i b_i > 0$. No other unresolved vintage can claim these deposits. Every effective discovery establishes that slot's false origin, decertifies the corresponding live unit, and forfeits these bonds. Innocent buyers have no liability for the organizer's fine. Side transfers deliver their face value, and each necessary mint participant accepts a payment equal to its expected incremental bond loss. The organizer has separate advance funding $F \geq \sum_i b_i$, so every minimum bribe schedule considered below is liquid.

The gateway charges a fixed fee ϕ per admitted unit. This finances ordinary verification wages, records, retained adjudication capacity and capital opportunity costs throughout the horizon. Honest wages cover each participant's inspection effort, bond funding and reservation payoff; there are no unfunded promised indemnities. The fee is prepaid and deposits are segregated. The batch invoice is therefore $h + \phi$; the buyer deposits it in escrow when accepting the contract. The fee funds service in advance, and the commodity component is a guaranteed payment to the supplier on admission, allowing clean procurement against the escrow without using the organizer's separate bribe budget. Buyers additionally finance their own investigation and retained downstream losses. They have enough liquidity for the invoice and the worst contractual loss. Bond forfeitures accrue to a public collector and are not needed to balance the gateway budget.

The organizer receives the commodity component h and either procures the admissible unit for h or supplies its existing illegal unit; ϕ passes to the gateway. All ordinary service compensation is unchanged by that choice. The stipulated inspection procedure is performed before the mint report, so collusion changes the report without saving its inspection cost.

Buyers observe the country and contract, but neither c nor the choice of backing. All admitted labels in the cohort pool. The clearing arrangement prevents buyer side contracts before evidence is obtained, including agreements to suppress effort. Once a buyer has verifiable evidence, it can approach the organizer and every original signer to bargain over silence before reporting to the authority. The finding itself does not notify the authority. A silence contract is enforceable, and the original participants can cofinance it and redistribute their gains. Thus the institution gives early evidence to a purchaser

who can choose how to use it. An independent late-enforcement channel remains beyond suppression. Collectible liabilities, continuation rents, and privately appropriable rewards consist solely of the stated bonds and buyer losses.

The stages are as follows.

1. The contract, fees, bonds and allocation rule are posted; buyers and honest service providers accept their participation contracts, buyers escrow their invoices, and each liable agent posts its own funded bond.
2. Each organizer privately observes c , chooses its backing and offers any necessary mint bribes. Inspectors otherwise reject illegal backing.
3. Admitted units are allocated and escrowed commodity payments are released. A buyer with posterior fraud probability π chooses effort $e \in [0, 1]$, costing $\kappa e^2/2$. Conditional on fraud, effort produces actionable early evidence with probability e , and never produces a false accusation.
4. A buyer obtaining evidence can bargain with the original coalition over a binding silence agreement, or report. Reporting triggers early decertification and net private loss $L_E \geq 0$. Without an early report, independent late enforcement finds the fraud with probability $r \in (0, 1)$ and produces net loss $L_L > 0$. Undiscovered fraud produces no additional private loss within this horizon.
5. The authority adjudicates, changes certification status, collects sanctions, and closes the exposure pool before deposits can be reused.

Early and late losses are incremental to the paid invoice and include retained claims work and operational interruption. Timely documented disclosure prevents an escalation of unresolved contractual exposure. The informed-silence continuation loss rL_L is already minimized over all feasible private mitigation, including quarantine or replacement. In this subclass, attaining the early-report loss requires documented disclosure. Silently stopping use cannot remove the earlier contractual exposure. These contractual restrictions give reporting its potential private advantage; no refund or reward is promised. Early exclusion can therefore impose a loss on an innocent buyer and still be preferable to silence. For the resource comparison, the losses consist of legal, replacement-logistics, and lost-service costs, with the original invoice and any damages transfer excluded from repeated counting.

4.11.2 Joint solution and disclosure incentives

Define $A = rL_L - L_E$ and impose the strict condition

$$A > (1 - r)D. \quad (71)$$

After learning with certainty that its unit is fraudulent, the buyer loses L_E by reporting and expects rL_L by remaining silent. It therefore requires a delivered payment of at least

A to accept silence. The entire original coalition together can save at most $(1 - r)D$: reporting forfeits its bonds with certainty, whereas silence leaves expected forfeitures rD . Its maximum willingness to pay is smaller than the buyer's reservation payment. Thus no silence agreement can leave all its participants weakly better off, even with coalition-wide cofinancing. Bilateral willingness to pay cannot exceed this coalition bound. Earlier mint bribes are sunk at this subgame; reshuffling them creates no extra suppression surplus.

The buyer reports every genuine finding, including an off-path finding. A buyer without verifiable evidence cannot establish a fraud claim or bargain over silence. Disclosure follows here from the joint surplus available in the evidence-stage bargain. With no equilibrium silence payment, expected additional buyer cost before learning is

$$R(\pi, e) = \frac{\kappa e^2}{2} + \pi\{eL_E + (1 - e)rL_L\} = \pi rL_L - \pi Ae + \frac{\kappa e^2}{2}. \quad (72)$$

Its unique best response is $e = \min\{A\pi/\kappa, 1\}$. Conditional on fraud, effective discovery has probability

$$p(e) = e + (1 - e)r = r + (1 - r)e. \quad (73)$$

This uses independence of late enforcement and the early signal, and counts only the first effective discovery.

At effort e , signer i requires advance compensation $p(e)b_i$ and the organizer bears its own expected bond loss $p(e)B_0$. The minimum funded capture cost is therefore $p(e)D$. Fraud is profitable for $c < \Gamma - p(e)D$; the remaining organizers procure clean material. At equality they choose honesty, a zero-mass tie under the continuous cost distribution. Buyers observe neither costs nor backing choices. Bayes' rule therefore makes their posterior fraud probability equal to the cohort's fraction of fraudulent admissions. Each buyer and organizer is negligible, so an individual deviation leaves that probability unchanged.

Theorem 4.9 (Positive-trade interior equilibrium). *Maintain (71), let $a = A/\kappa$, $T = \Gamma - rD$, and assume*

$$0 < T < C, \quad a > 0, \quad aT < C + D(1 - r)a. \quad (74)$$

Subject to buyer participation below, the posted-contract market has a unique fraud fraction and audit choice:

$$\pi^* = \frac{T}{C + D(1 - r)a}, \quad e^* = a\pi^*, \quad p^* = r + (1 - r)e^*. \quad (75)$$

All one unit of slots trade; $0 < \pi^ < 1$, $0 < e^* < 1$, original signers accept their minimum compensating bribes when offered, and buyers disclose all actionable evidence. Honest suppliers and verifiers participate, and both initial bribes and continuing enforcement obligations are funded.*

Proof. Strict convexity in (72) gives the buyer's unique best response. For an interior solution the uniform cost distribution gives $C\pi = \Gamma - rD - D(1 - r)a\pi$, yielding (75).

The solution lies strictly inside the unit interval because its numerator is positive and less than C ; the last inequality in (74) ensures $e^* = a\pi^* < 1$. More generally, the fraud best response is the continuous, nonincreasing map $\pi \mapsto [\{\Gamma - Dp(\min(a\pi, 1))\}/C]_0^1$, where brackets truncate to $[0, 1]$. Its difference from π is strictly decreasing; hence another fixed point is impossible. Each signer is indifferent between its funded offer and its maintained honest payoff, and $p \sum_i b_i \leq F$. Sequential reporting and the absence of mutually acceptable silence contracts were established above. Posted wages, prepaid fees and segregated deposits satisfy the stated participation and budget contracts; buyer participation completes the positive-trade allocation. $\square$

If each buyer values the underlying contracted service at v , its expected net payoff is $v - h - \phi - R(\pi^*, e^*)$. Thus the sufficient, explicit participation condition is

$$v > h + \phi + R(\pi^*, e^*). \quad (76)$$

One buyer is contracted to each slot, making accepted demand and admitted deliveries both equal to one. The participation inequality sustains this quantity at the posted invoice. If it fails, the positive-trade allocation is unavailable. Endogenous competitive prices and aggregate demand are introduced when the certification institution is embedded in the production economy.

4.11.3 The effect of downstream scrutiny

The matched mint-only benchmark keeps h, p_D, C, D, r, L_L and the original participants, but omits the downstream evidence channel: $e_0 = 0$ and $\pi_0 = T/C$. Giving that benchmark the same late decertification rule makes the comparison depend on additional scrutiny and its deterrent effect. Then $\pi^* < \pi_0$ and $p^* > r$. With no false-positive removals or exits, live contamination is

$$\iota(\pi, p) = \frac{\pi(1-p)}{1-\pi p}. \quad (77)$$

Its derivatives are $(1-p)/(1-\pi p)^2 > 0$ with respect to π and $-\pi(1-\pi)/(1-\pi p)^2 < 0$ with respect to p in this interior class. Therefore $\iota(\pi^*, p^*) < \iota(\pi_0, r)$, accounting for the smaller surviving denominator as well as fewer illegal admissions.

The change in certification redirects physical stocks. Clean procurement is $1 - \pi$ and diversion is π . The original unit of dirty stock supplies $1 - \pi$ to the outside continuation. Among admitted units, πp become excluded inventory and $1 - \pi p$ remain certified, including $\pi(1-p)$ illegal mass. Certified, excluded, and outside-delivered quantities sum to $2 - \pi$: clean production $1 - \pi$ plus the original dirty stock. Greater credibility in this example raises dirty outside deliveries by reducing diversion. Illegal extraction is fixed. At the same time, the organizer's expected incremental rent over faithful conduct is

$$\int_0^{C\pi^*} \frac{\Gamma - p^*D - c}{C} dc = \frac{C(\pi^*)^2}{2}. \quad (78)$$

The equality uses $\Gamma - p^*D = C\pi^*$. Expected forgery resource expenditure is separately $\int_0^{C\pi^*} c dc/C = C(\pi^*)^2/2$.

Let the baseline fee be ϕ_0 and let maintaining the new channel require additional fixed record, capacity and exposure costs $H > 0$, so $\phi_1 = \phi_0 + H$. Recurring prevention costs rise by $H + \kappa(e^*)^2/2$. Expected incident losses can fall, since fewer frauds occur. Total institution and downstream resource costs are

$$\mathcal{C}_0 = \phi_0 + \pi_0 r L_L, \quad \mathcal{C}_1 = \phi_0 + H + R(\pi^*, e^*). \quad (79)$$

Fees in this account equal service and capital opportunity costs. Adding fees to the wages they finance would count the same expenditure twice; bond seizures and bribe receipts are also transfers. Total institutional cost rises precisely when $H > \pi_0 r L_L - R(\pi^*, e^*)$. Positive diligence alone does not determine the sign because fewer frauds reduce incident losses. Commodity production, forgery resources, and the dirty-market opportunity value enter the wider mineral-economy account separately.

Example 4.10 (Lower leakage, higher funded cost). Take $h = 3$, $p_D = 2$, $C = 1$, $r = 1/4$, $B_0 = 1$, and two distinct required participants, a licensed custodian and an inspector, with bonds $1/2$ each. Thus $D = 2$ and $F = 1$ suffices. Set $L_E = 1$, $L_L = 12$, $\kappa = 2$, $\phi_0 = 1$, $H = 1$, and $v = 7$. Then $A = 2$, $a = 1$ and $T = 1/2$ satisfy the theorem's strict conditions, including $A = 2 > 3/2 = (1 - r)D$. The exact outcomes are

	Mint-only benchmark	Continuing scrutiny
Fraudulent admissions π	$1/2$	$1/5$
Buyer effort e	0	$1/5$
Conditional discovery p	$1/4$	$2/5$
Live certified mass	$7/8$	$23/25$
Live contamination	$3/7$	$3/23$
Buyer additional cost R	$3/2$	$14/25$
Institutional resource cost $\mathcal{C}$	$5/2$	$64/25$
Invoice plus expected buyer cost	$11/2$	$139/25$

The new invoice is 5, and buyer expected surplus is $36/25 > 0$. The new minimum total mint bribe is $2/5$, below available advance funding; each signer receives $1/5$, covering its expected forfeiture. The organizer also expects its own forfeiture of $2/5$, so precisely costs $c < 1/5$ yield profitable fraud. All deposits remain collectible through both discovery dates. After evidence, the buyer requires at least 2 for silence, while the entire original coalition can save at most $3/2$ in sanctions. Added recurring prevention costs are $26/25$, while incident losses fall from $3/2$ to $13/25$, a saving of $49/50$; their net increase is $3/50$. These are illustrative exact parameters, not estimated leakage or cost magnitudes.

True origins determine where this fraud margin is available. For known countries outside $\mathcal{I}$, its action set is empty: $\pi = 0$ and investigation motivated solely by this fraud is zero. Ordinary provenance costs can remain. When exposed and safe origins trade together, the live-mass weights in Section 4.1 determine aggregate leakage. The country composition of actual supply consequently matters for both remaining exposure and the incidence of prevention costs.


4.12 From credibility to production incentives

Certification affects the rest of the economy through the mineral it redirects and the returns it leaves to market participants. For each true origin, physical accounting distinguishes extraction, diversion from the unverified outlet, displaced legal production, certified deliveries, subsequent withdrawals, and physical losses. A withdrawn claim can lead to storage, lawful outside sale, replacement, or destruction; each has its own physical consequence. Fraudulent admission changes eligibility without changing the legal status of the original extraction.

The corresponding producer incentive is the certified receipt net of fees, operating costs, and continuing liabilities. The buyer compares that contract with its full purchase and diligence costs. The organizer also gives up the value of the dirty stock it diverts. With an auction or resale continuation, this loss is evaluated under the committed mechanism and information, retaining the original stock and menus in $W_{\Sigma}(q - d \mid q, \mathcal{H})$. The posted-contract market above uses an explicit linear outside quote; the subsequent equilibrium analysis determines receipts and opportunity costs jointly with production.

These opposing forces explain why improved credibility need not have a uniform effect on rents. Greater confidence can support demand for certified supply, while the resources used to establish it reduce the net return to that supply. Less diversion can leave more dirty mineral for an outside market. The effects on clean investment, illegal extraction, and government revenue therefore depend on market clearing as well as on the certification institution. Persistent exposure makes false certification harder to sustain under the conditions established here. Financing that exposure is part of the economic adjustment through which provenance can change production at its origin.

5 Regional production and world supply



Regional supply enters world markets, while receipts feed back through regional allocations.

A policy that changes the return to mineral production changes more than the quantity sold abroad. It changes which regions produce, the terms on which households work, and the income available to them. To follow these responses through the world market, aggregation must preserve the regional choices that generate supply. The same construction must allow a world allocation to be traced back to its producing households.

For the separable production technology considered here, this can be done exactly. Regions with greater productive capacity or lower implementation cost supply more at a common receipt. Their contributions add to a single capacity measure, while the underlying regional weights determine the allocation of any aggregate quantity. Geography enters through these weights and the mass of regions, rather than through an assumed representative producer.

There are finitely many origin, technology and contract groups g . Group g contains an atomless population $(\mathcal{R}_g, \mu_g^{\text{reg}})$ of regions, with mass $m_g = \mu_g^{\text{reg}}(\mathcal{R}_g)$. Each region has the attained household implementation correspondence $\Gamma_r(q; A_r)$. On the installed-equipment branch of the Selva economy, its private cost is

$$\mathcal{C}_r(q; A_r) = \frac{\eta}{1+\eta} \kappa_r A_r^{-1/\eta} q^{1+1/\eta}, \quad \beta_r = A_r \kappa_r^{-\eta}, \quad B_g = \int_{\mathcal{R}_g} \beta_r d\mu_g^{\text{reg}} = m_g \mathbb{E}_g[\beta_r]. \quad (80)$$

The elasticity $\eta > 0$ is common, $0 < B_g < \infty$, and zero output is feasible. Implementations are separable across regions, with the measurability and resource conditions of Assumption C.1. The private cost includes both productive inputs and the operating income that must remain with household owners to induce the chosen allocation.

First suppose producers in a group receive a common amount per unit: $W_r(q; \Omega) = r_g(\Omega)q$. The receipt $r_g(\Omega)$ is net of the contract's variable payments and liabilities, and each region takes the public state Ω as given. A constant receipt is thus a property of the trading

institution. Where fraud or retail market power makes the continuation nonlinear, its entire value remains in the regional problem $\max_q \{W_r(q; \Omega) - C_r(q; A_r)\}$.

Proposition 5.1 (Exact geographic aggregation). *Under these conditions, output, operator profit and the least private cost of implementing group output are*

$$\begin{aligned} Q_g &= B_g [r_g(\Omega)]_+^\eta, & \Pi_g &= \frac{B_g [r_g(\Omega)]_+^{1+\eta}}{1+\eta}, \\ \bar{C}_g(Q) &= \frac{\eta}{1+\eta} B_g^{-1/\eta} Q^{1+1/\eta}, & q_r(Q) &= \frac{\beta_r}{B_g} Q. \end{aligned} \quad (81)$$

The regional reconstruction is measurable, integrable and unique almost everywhere among allocations attaining the aggregate cost. If the regional household correspondences are measurably selectable and integrably bounded at these outputs, it admits a feasible measurable, integrable household reconstruction. The constructive installed-equipment class below satisfies these conditions under the stated bounds.

Proof. At positive receipt, the strictly concave regional objective has first-order condition $r_g = \beta_r^{-1/\eta} q_r^{1/\eta}$; at nonpositive receipt its unique optimum is zero. Integrating gives output and profit. For a fixed total Q , normalize the measure to $\beta_r d\mu_g^{\text{reg}}/B_g$ and apply strict convexity to $(q_r/\beta_r)^{1+1/\eta}$. The minimum occurs exactly when $q_r/\beta_r = Q/B_g$. This explicit output allocation is measurable and has integral Q . The selection argument and complete household construction are proved in Appendix C, Propositions C.6 and C.7. $\square$

The allocation equalizes marginal implementation cost. A region's share of output is its share of effective capacity, so an increase in group production expands each region in the same proportion. Doubling the mass of identical regions doubles B_g , output and profit; subdividing that population leaves all three unchanged. This gives world supply an exact regional interpretation.

The cost envelope also applies to a coordinated purchaser whose downstream value depends only on total Q : conditional on that total, it allocates procurement to minimize cost. Local sellers with separate nonlinear continuations face different marginal revenue schedules. Their production choices must instead be solved locally and then integrated. Two such populations can have the same B_g and different supply because capacity is distributed differently across local markets.

5.1 Who produces and who receives the rents

Groups with the same physical origin j may receive different shares of the world price. Let the gross price be $p_j \geq 0$ and let the nonnegative multiplicative receipt be $r_g = \omega_g p_j$. Keep certification, access and contract groups separate. Output and profit are then

$$Q_j = p_j^\eta \sum_{g \in j} B_g \omega_g^\eta, \quad \Pi_j = \frac{p_j^{1+\eta}}{1+\eta} \sum_{g \in j} B_g \omega_g^{1+\eta}. \quad (82)$$

A higher receipt share encourages production and increases the return on each unit produced. The supply coefficient captures the first effect; the profit coefficient captures both. This accounts for the extra power of ω_g in the second expression. A government valuing protection income, or an owner valuing an invention, therefore needs the profit coefficient together with its own fiscal or ownership claim. Supply alone does not reveal the income over which these actors compete.

Additive fees, output-dependent liabilities and optional contracts change the regional optimization. In those cases aggregation proceeds by integrating the resulting choices directly. The two moments in (82) characterize the multiplicative-receipt branch.

5.2 Household income behind aggregate supply

The return to households can be seen explicitly on the constructive Selva branch. Put $\alpha_m = \eta/(1 + \eta)$. A household-owned firm produces $q = a_r g^{\alpha_m}$, necessities sell at resource prices, and the input price must be at least its resource cost $c_r^g > 0$. At a positive target q_r , the cost-minimizing input quote is c_r^g and the mineral procurement quote is $\theta_r = \mathcal{C}_r(q_r)/q_r$. The revenue account is

$$c_r^g g_r = \alpha_m \mathcal{C}_r(q_r), \quad H_r = (1 - \alpha_m) \mathcal{C}_r(q_r), \quad r_g q_r = \Pi_r + c_r^g g_r + H_r. \quad (83)$$


Here H_r is the operating income of household owners. Nonmining endowments finance a strictly surviving baseline; each household then allocates its ownership share of H_r between necessities and cash. Proposition C.7 gives these choices, including households whose optimal cash holding is zero. The same procurement payment finances the firm's input purchases and its owners' income.

At the common receipt, $\theta_r = \alpha_m [r_g]_+$ and output is proportional to β_r . Identical supply responses nevertheless allow different distributions of welfare. Ownership, outside wealth and subsistence needs determine how the receipt translates into consumption. Likewise, environmental consequences depend on the extraction, input and site-use activities in Γ_r . The capacity measure B_g determines supply on this branch; household and environmental outcomes require their corresponding allocations.

Fixed operating costs introduce an additional margin. At low receipts, regions exit; at a participation threshold, an atomless population can support any active fraction. The supporting-receipt construction in Proposition C.9 gives an attained cost envelope even with heterogeneous fixed costs. Capacity ceilings, subsistence output floors and linked village coverage require the constrained choices in (156). Each branch carries its own feasible household allocation. In particular, exit is an economic option only where household survival outside production is financed. The general regional objects remain $\mathcal{C}_r$ and Γ_r throughout these comparisons.

This regional foundation makes it possible to ask how certification changes production. The answer requires the prices producers actually receive, the quantities admitted to each market, and the cost of sustaining the claims on which buyers rely.

6 Certification and world-market equilibrium



Certification separates markets; prices transmit its effects to regional production.

6.1 The contract buyers purchase

A certificate can change where a mineral is sold and what a buyer will pay for it. Its effect on production depends on the resulting return at the source. When illegal material retains an outside outlet, better screening may change the composition of certified deliveries without reducing total illegal extraction. We therefore determine certification, procurement and world prices jointly, starting with a market in which suppliers receive public per-unit quotes.

The contract supplies one physical unit together with its traceable origin record. Before purchase, the buyer accepts a public lottery over the origins actually supplying the market. After assignment, it observes the origin and its contract and chooses diligence accordingly. Each origin keeps its own history and evidence. Clean suppliers receive the common certified quote $h > 0$. Let $x = p_D/h$ be the dirty-to-certified supplier-price ratio, so dirty suppliers receive $p_D = xh$. A separate licensed intermediary organizes admission to the certified market and purchases any diverted dirty material at p_D . This intermediary owns the return to its fraud. Common ownership with a cartel would require consolidating that return and solving the owner's changed production decision.

There are finitely many true origins k , each containing the regional populations of Section 5. Let B_{Ck} denote clean supply coefficients and $B_D = \sum_k B_{Dk}$. An origin is vulnerable when illegal physical supply and an admissible production branch coexist. An origin with no illegal stock has no illegal-origin admission action. At a vulnerable slot, the forgery-cost draw is uniform on $[0, c_k]$; its monetary cost is h times the draw. Bonds, fees, incident losses and diligence costs also scale with the commodity invoice h . Sustaining this contract as physical volume changes requires paid verification productivity and the funding conditions developed in Section 8.

6.2 The incentives to investigate and report

Verification uses resources, and information changes behavior only when those who obtain it have an incentive to disclose it. The first issue is central to costly state verification [Townsend, 1979]. The second has direct empirical relevance: changing auditors' incentives in Gujarat improved reporting and reduced plant emissions [Duflo et al., 2013]. Here both investigation and disclosure are individual choices, and the gains from false admission depend on the endogenous price difference between markets.

Let $r_k \in (0, 1)$ be the probability of unsuppressible late discovery, b_k the sum of distinct collectible bonds per unit of the commodity invoice, so their monetary total is hb_k , and hl_{Ek} and hl_{Lk} the buyer's losses after an informed early report and conditional on later discovery following silence, respectively. The expected loss from informed silence, after all feasible private mitigation, is $hr_k l_{Lk}$. Diligence costs $h\kappa_k e^2/2$. Assume $c_k, \kappa_k > 0$ and $l_{Ek} \geq 0$, and define

$$a_k = \frac{r_k l_{Lk} - l_{Ek}}{\kappa_k}, \quad T_k(x) = 1 - x - r_k b_k, \quad D_k = c_k + b_k(1 - r_k)a_k. \quad (84)$$

Maintain $r_k l_{Lk} - l_{Ek} > (1 - r_k)b_k$, separate advance finance for bribes, and the information and reporting access specified in the certification game. A buyer with evidence then requires more compensation for silence than the original coalition can jointly save. Reporting is optimal. On the interior contract branch, fraud and diligence satisfy

$$\pi_k(x) = \frac{T_k(x)}{D_k}, \quad e_k = a_k \pi_k, \quad p_k = r_k + (1 - r_k)e_k, \quad R_k = \pi_k r_k l_{Lk} - \frac{(r_k l_{Lk} - l_{Ek})^2 \pi_k^2}{2\kappa_k}. \quad (85)$$

The variable p_k is a discovery probability; the dirty-market price is p_D . A safe origin has $\pi_k = e_k = R_k = 0$. The clipped fraud fixed point is unique, and the smooth price analysis uses vulnerable origins with $0 < \pi_k < 1$ and $0 < e_k < 1$.

The equations display the two forces governing false admission. A larger clean–dirty price difference raises its private gain. Anticipated investigation and collectible sanctions lower that gain. Buyers in turn investigate more when fraudulent admission is more likely. The market price ratio determines where these decisions meet.

6.3 Market composition and the destination of production

Regional optimization gives $Q_{Ck} = B_{Ck}h^\eta$ and $Q_{Dk} = B_{Dk}(xh)^\eta$. Honest procurement supplies the fraction $1 - \pi_k$ of origin k 's admissions. Hence

$$X_k = \frac{B_{Ck}h^\eta}{1 - \pi_k}, \quad D_k^{\text{div}} = \pi_k X_k, \quad Y_k = Q_{Dk} - D_k^{\text{div}}. \quad (86)$$

Impose $Y_k > 0$ at every vulnerable origin: false admission must use material available at that source. Let $X = \sum_k X_k$ and $w_k = X_k/X$ be the actual composition of admissions. Buyers' expected diligence and retained losses are then $h \sum_k w_k R_k$. With a common

funded service fee $h\phi$, the generalized unit cost of the standardized contract is $hK(x) = h[1 + \phi + \sum_k w_k R_k]$. Composition matters because deliveries from different origins carry different expected verification costs and losses.

Physical production is exhausted by certified admissions and outside deliveries:

$$\sum_k (Q_{Ck} + Q_{Dk}) = X + Y, \quad Y = \sum_k Y_k. \quad (87)$$

Subsequent revocation divides admissions into live mass $\sum_k X_k(1 - \pi_k p_k)$ and excluded mass $\sum_k X_k \pi_k p_k$. Both remain physical material. In this economy delivered mineral is embedded in serialized durable products, with no recycling or return to current primary supply. Existing material and permanent evidence stay in the state. Product resale would transfer those holdings; recycling into new primary deliveries would require an explicit release equation. This physical account separates changes in legal status from new extraction.

6.4 Prices, production and admission

Final purchasers have outside-good endowments and quasilinear preferences $E_C \log X + E_D \log Y + n$, where $E_C, E_D > 0$. They take prices and the public origin lottery as given, can finance their contracts and maximum specified losses, and choose diligence after assignment. Demand satisfies $hKX = E_C$ and $xhY = E_D$. The first expression includes both the invoice $h(1 + \phi)X$ and the separately borne expected costs $h \sum_k X_k R_k$. Only the invoice is received by the supplier and service providers.

Define the admission coefficient F_{cw} , the diversion coefficient H_{cw} , the certified-expenditure coefficient T , and the expenditure-ratio function J by

$$F_{cw}(x) = \sum_k \frac{B_{Ck}}{1 - \pi_k(x)}, \quad H_{cw}(x) = F_{cw}(x) - \sum_k B_{Ck}, \quad (88)$$

$$T(x) = \sum_k \frac{B_{Ck}[1 + \phi + R_k(x)]}{1 - \pi_k(x)}, \quad J(x) = \frac{x[B_D x^\eta - H_{cw}(x)]}{T(x)}. \quad (89)$$

Take a closed interval $[\underline{x}, \bar{x}] \subset (0, 1)$ throughout which the interior incentive conditions and all origin-specific stock inequalities hold.

Theorem 6.1 (Joint common-price equilibrium). *Suppose the individual incentive, resource and liquidity conditions above hold and $J(\underline{x}) < E_D/E_C < J(\bar{x})$. There is a unique equilibrium in this contract branch. It is characterized by*

$$J(x^*) = E_D/E_C, \quad h^* = \left[\frac{E_C}{T(x^*)} \right]^{1/(1+\eta)}, \quad p_D^* = x^* h^*. \quad (90)$$

All output, composition, fraud, effort and live claims follow (85)–(86). Every regional output is implementable by its household allocation correspondence. On compact sets of


strictly feasible primitives the equilibrium is continuously differentiable, with uniformly bounded derivatives.

Proof. Each π_k decreases with x . On the interior effort branch, $dR_k/d\pi_k = r_k l_{Lk} - (r_k l_{Lk} - l_{Ek})e_k \geq l_{Ek} \geq 0$. Thus H_{cw} decreases and T decreases. Where residual dirty delivery is positive, the numerator of J strictly increases. Therefore $J' > 0$. The endpoint inequalities, continuity and intermediate value theorem give exactly one x^* ; clean clearing gives h^* . Substitution proves both markets clear and all physical constraints hold. The individual fraud and audit fixed point and the regional maximizing rules are already verified. Compact strict margins bound J' away from zero, so the implicit function theorem gives the last assertion. Full reconstruction and the cash accounts are in Appendix D. $\square$

At a fixed clean quote, an increase in the dirty price expands dirty production and reduces the gain from diverting it into certification. More material therefore remains for the outside market. Lower fraud also reduces each vulnerable origin's invoice-normalized diligence and retained losses, and the certified-expenditure coefficient $T(x)$ decreases. These comparisons make $J(x)$ strictly increasing in the price ratio; clean-market clearing then determines the level of h . The equilibrium composition of trade follows from the same incentives that determine supply.

Example 6.2 (An equilibrium with endogenous production and verification). For one vulnerable origin take $\eta = 1$, $B_C = 4/15$, $B_D = 1/2$, $r = 1/4$, $b = 2/3$, $c = 1/3$, $l_E = 1/3$, $l_L = 4$, $\kappa = 2/3$, $\phi = 2/3$, $E_C = 139/25$ and $E_D = 8/5$. The unique interior clearing solution is $h = 3$, $p_D = 2$, $\pi = e = 1/5$, $p = 2/5$, $X = 1$, $Q_C = 4/5$, $Q_D = 1$ and $Y = 4/5$. The live contamination rate is $3/23$. These quantities follow from regional profit maximization and buyer demand. The displayed coefficients can be split into unequal regions without changing the equilibrium, while preserving every regional output.

A. Physical deliveries



B. Status of the admitted cohort after scrutiny




Figure 4: Deliveries and claim status in Example 6.2. Panel A conserves mineral: total extraction and total delivery both equal $9/5$; panel B follows the admitted unit through scrutiny, with $23/25 + 2/25 = 1$.

6.5 Production rents and surviving contamination

The returns to clean and dirty producers are

$$\Pi_C = \frac{h^{1+\eta}}{1+\eta} \sum_k B_{Ck}, \quad \Pi_D = \frac{(xh)^{1+\eta}}{1+\eta} \sum_k B_{Dk}. \quad (91)$$

Licensed organizers receive an additional $\sum_k hc_k \pi_k^2 X_k / 2$ above their compensation for honest service. This is their profit after the costs and expected sanctions of fraud. The uniform-cost specification gives the same numerical expression for the real forgery bill, but the recipients and economic meanings differ. Fees, bribes, seizures and commodity payments transfer income; forgery, investigation and productive inputs consume resources. Household income inside the implementation cost remains part of the distribution of receipts. These distinctions determine which returns can finance innovation or attract government protection.

For dated cohorts, the share of illegal material among surviving live claims is


$$l_t^{\text{live}} = \frac{\sum_{k,v} \mathbb{E} U_{kvt}}{\sum_{k,v} \mathbb{E} V_{kvt}}, \quad (92)$$

provided the denominator is positive. The numerator and denominator follow the same physical cohorts through time. A reform changes the status of new production from the next date; it leaves the extraction origin of old material unchanged.

The construction in (164) makes surviving contamination arbitrarily small and strictly positive within this interior class by raising the real forgery-cost scale, maintaining funded reporting, and choosing expenditure ratios within the strict price interval. The production consequence still depends on p_D and regional productive capacity. Illegal suppliers are paid for material entering either destination. Screening success and lower illegal extraction are therefore separate outcomes whose connection must be established through equilibrium prices.

The public outside-sale route used here is committed before production and gives $W(q) = p_D q$. When access to processing is exclusive, the seller instead faces a nonlinear return to its stock. The next section carries that selling problem into world equilibrium and examines how its allocation affects prices and rents.

7 Retail market power and the world price



Exclusive processing routes connect local selling decisions to final demand and production receipts.

A public mineral price does not eliminate local market power when access to processing remains scarce. A seller that controls a particular route to downstream buyers chooses how much to supply through that route and how much to retain. World demand then values the final output created by those choices. To understand the effect of reform or innovation on this seller, we must carry its allocation and payment rules into market clearing.

We compare two selling institutions. With a continuum of unit-demand buyers, the Muir mechanism admits an equilibrium with a positive rationing premium, resale and retained conversion. With finite whole buyers, the VCG mechanism has a nonlinear return to inventory and an efficient complete-game equilibrium with positive expected retention. In each case, regional production responds to the seller's optimized return under its specified buyer population and commitment assumptions.

7.1 Access to processing and the return to production

A region contains an atomless mass n_r of independently operated production and retail cells. Each cell is negligible in the world market but serves the buyer population specified by its selling institution. A processor with productivity v converts one unit of mineral into v units of a final product sold competitively at price p . Its direct mineral value is pv .

Processors have exclusive, enforceable access to their cell's mineral. They cannot obtain an equivalent publicly traded input, create additional identities, or enter another cell. Producers likewise have geographic or enforceable origin-specific access only to their assigned processors and the retained conversion route. This restriction is fixed before participation and procurement. Their alternative is exit at zero; declining the contract does not open the bulk raw-mineral outlet. The custody arrangement below deters covered unauthorized route changes. The shared final-product market therefore connects demand while preserving the restrictions on raw-input access that support local market power.

A public consignment contract divides regional stock equally across installed cells at every chosen regional quantity. With the regional power-cost envelope, the resulting continuation is $W_r(Q; p) = n_r W_j(Q/n_r; p)$. Each cell retains its own buyers and mechanism. This routing rule is the link between the regional procurement problem and downstream selling.

Unsold mineral can be converted into the same final product by a separate technology $H(y)$. Retention yields $pH(y)$, and the producer cannot subsequently release that stock to retail buyers. Every extracted unit enters either primary delivery or the retained route. Resale reallocates mineral already delivered.

For a Muir cell, let $V(m)$ be optimal primary receipts at price one, including the repeated-acquisition game and its competitive resale selection. Direct values and the specified acquisition charge both scale with p . Production and retention solve

$$\pi(p) = \max_{m, y \geq 0} \left\{ p[V(m) + H(y)] - \frac{k}{2}(m + y)^2 \right\}, \quad q = m + y. \quad (93)$$

Equivalently, the seller chooses q to maximize $W(q; p) - \mathcal{C}(q)$, where $W(q; p) = p \max_{m \leq q} \{V(m) + H(q - m)\}$. Procurement therefore anticipates both the premium available from primary selling and the opportunity value of retained mineral.

Let $Z(p)$ be final output from the actual primary allocation, resale and retained conversion. This physical output differs in general from receipts divided by p . Processing rents,

acquisition charges, household income and producer profit divide the value created by conversion among different actors. Across finitely many origin and contract types j , final supply is $Z_R(p) = \sum_j n_j Z_j(p)$. Exact atomless replication gives the specified aggregate frequencies of private types and resale outcomes; one cell has no effect on the world price.

A final consumer with sufficient numeraire has utility $c_0 + E_R \log Z_R$ and demand E_R/p . The world price solves

$$p \sum_j n_j Z_j(p) = E_R. \quad (94)$$

Thus market clearing uses the allocation generated by local selling. Payments determine the incentive to produce, while final physical output determines what consumers can buy.


Proposition 7.1 (An exact active nonlinear class). *There is a nonempty open parameter class containing the explicit Muir construction in Appendix E in which each region has a unique global production choice, strictly positive retained deliveries, an active rationing premium, and strictly positive resale. On a compact price interval its $q(p)$, $Z(p)$ and actual profit $\pi(p)$ are smooth; $q'(p) > 0$, $Z'(p) > 0$ and $\pi'(p) > 0$.*

The separate finite VCG construction in that appendix also has an attained global production choice and positive expected retained mineral. Its inventory continuation changes at whole-package capacities. On its strict branch, physical output is constant and actual profit is increasing in p ; no active secondary trade is claimed for its efficient equilibrium.

For any collection of these cells, if the two endpoint signs in (94) enclose a root, the root in that interval is unique and continuously depends on type masses and demand. Actual regional production and profit are $q_r = n_r q_j(p)$ and $\Pi_r = n_r \pi_j(p)$, and their cost-minimal household implementations are attained. This operator remains continuous when one type's mass tends to zero, provided a positive core remains.

The Muir seller expands production as the final-product price increases, while balancing primary receipts against retained conversion. The finite VCG seller faces a different margin: the number of whole packages it can offer. Over the strict price interval of the construction, its preferred stock stays fixed even as its expected profit rises. Both responses enter aggregate supply through the actual cell solution.

Physical output from one assigned retail origin



Selling branch	Physical response to higher p	Producer return
Muir	$q'(p) > 0$ and $Z'(p) > 0$	Actual profit rises; the stated branch has active resale.
Finite VCG	Preferred inventory stays fixed on its strict branch	Expected profit rises; no active secondary trade is claimed.

Figure 5: Physical output and production incentives under exclusive retail access. Each selling institution generates its own allocation and payments; the branch comparisons are those of Proposition 7.1.

Replication connects these choices to heterogeneous regions. For the quadratic cell cost $kq^2/2$, regional least cost is $kQ_r^2/(2n_r)$ and hence $\beta_r = n_r/k$. More generally, a cell cost $k\eta y^{1+1/\eta}/(1+\eta)$ gives $\beta_r = n_r k^{-\eta}$. Regional size is therefore preserved in production, cost and profit. An innovation that replicates a cell must install its production capacity, secure its processing access and finance its institutions. An improvement in extraction efficiency inside an existing cell changes its optimization and must be valued through that changed problem.

7.2 Substitution across the certified, outside and retail markets

The same consumers can demand certified mineral, outside dirty mineral and the processed retail product. Give the common final consumer the strictly concave utility

$$c_0 + E_C \log X + E_D \log Y + E_R \log Z_R + \varepsilon \log(Y + \gamma Z_R), \quad \varepsilon, \gamma > 0. \quad (95)$$

Here X denotes the standardized certified contract, Y outside dirty delivery, and Z_R processed retail output. The last term introduces substitution between outside use and the processed product. Bulk and exclusive retail routes use disjoint origin and contract

groups, each with its own reconstructed regional production. Their combined extraction supplies the world market, with each mineral unit assigned to one route.

Using $p_D = xh$, $X = h^\eta F_{\text{cw}}(x)$ and $Y = h^\eta [B_D x^\eta - H_{\text{cw}}(x)]$, joint clearing requires

$$\begin{aligned} hK(x)X &= E_C, \\ hxY &= E_D + \frac{\varepsilon Y}{Y + \gamma Z_R(p)}, \\ pZ_R(p) &= E_R + \frac{\varepsilon \gamma Z_R(p)}{Y + \gamma Z_R(p)}. \end{aligned} \tag{96}$$

The price of processed output now influences demand for outside mineral, and the outside price influences retail demand. These responses change the gains to fraud and the returns to regional production. Certification composition, scrutiny and net procurement quotes continue to come from the bulk certification game; retail supply contributes its own production, payments and final output. Reform in either set of origins can therefore change the prices faced by the other.

Theorem 7.2 (Connected nonlinear and certification equilibrium). *Take a compact neighborhood of interior bulk equilibria with nonsingular price Jacobian, positive X, Y , strict institutional funding and incentive margins, and no clipping switch. Combine it with the strict retail branch of Proposition 7.1, with positive retail core and endpoint clearing margins. For all sufficiently small $\varepsilon > 0$, (96) has a continuously differentiable equilibrium branch on a sufficiently small compact neighborhood of the baseline state. Its Muir premium, resale, positive retained deliveries, actual funding and physical balances persist. Both price systems and the actual producer values respond to the common demand system. The assertion is local existence and uniqueness of this branch, not global uniqueness among every possible auction, certification or price regime.*

The equilibrium follows from the full price Jacobian and consumers' first-order conditions. For a continuing sequence of reforms, the state must also remain within the compact domain on which these conditions hold. In particular, a fixed positive outside expenditure requires positive outside delivery at bounded prices. A retail type can have zero aggregate mass while other retail suppliers remain active, but that is a different boundary from disappearance of the outside market.

Theorem E.4 establishes a growth path satisfying these restrictions. Safe bulk supply and retail installations expand at scale L_t , while vulnerable-origin capacity, outside dirty supply and outside expenditure remain in absolute units. The three-price system is uniformly smooth in $\lambda_t = 1/L_t$ through $\lambda = 0$. A fixed positive vulnerable capacity and a small positive substitution coefficient preserve endogenous low leakage and price feedback at every finite date. Outside dirty extraction retains its absolute positive core even as its share of the expanding economy declines.

Credible retail provenance requires control of the physical route as well as evidence about it. Proposition E.5 assumes that every covered false-origin or diversion action needs a knowing bonded participant who observes the relevant origin and action and can veto

it. Registered custody, accessible inspectors and processors, and an unsuppressible late audit make the exposed liabilities enforceable. Distinct collectible pools exceed every finite coalition gain under the committed continuation. Independent custody providers own these pools; they are separate from producer wealth.

A nonmineral treasury appropriation pays for installation, capital carry, service, evidence storage and segregated reserves. Within this institutional class, false retail admission is deterred. The separate vulnerable bulk origins retain their positive organizer profit and surviving contamination. Both the resources needed to maintain retail custody and the remaining bulk fraud enter the dynamic economy.


7.3 The income available to innovators and governments

Consider balanced replication, with $n_{j,t} = L_t \hat{n}_{j,t}$ and $E_{Rt} = L_t e_{Rt}$. Here $\hat{n}_{j,t} = n_{j,t}/L_t$ is type j 's cell mass per unit of scale and $\hat{n}_t = (\hat{n}_{j,t})_j$ is its vector. Dividing (94) by L_t leaves a price determined by $(\hat{n}_t, e_{Rt})$, and type profit is $L_t \hat{n}_{j,t} \pi_j(p_t)$. The same normalization applies to (96) when capacities, expenditures and the substitution coefficient all scale by L_t . The mixed-scale path above instead keeps its vulnerable and outside cores in absolute units. In either case, an individual research owner takes the public price path as given, while a positive mass of successful projects changes that path through clearing.

For an independently licensed additional cell, a one-period private lead can yield its dated operating profit $\pi_j(p_{t+1})$ less installation, access and institutional costs, multiplied by the number of new owned cells and the discount factor. This shorter prize applies when later funded opportunities and residual liabilities are identical after success and failure. If licensing income or exposure outlasts the lead, the remaining dated payments stay in the owner's return. Public disclosure can end exclusivity while leaving the financial obligations of the additional production intact.

The seller's prior commitment also governs diversion. Its return uses $\mathcal{W}_\Sigma(q - d \mid q, \Omega)$ at market state Ω , preserving the original menus, retained route, information, invoices and default rules. An optimized $W(q - d)$ is available only when a verified pre-report stock contingency was part of the original commitment. Procurement already paid is sunk in this comparison. These actual receipts and obligations determine the resources available for innovation and government intervention, to which we now turn.

8 Technical change and the incentive to govern



Funded investment and fiscal incentives can bring public authority into Selva, while funded livelihoods and safe production sites support protected-site conservation.

A more profitable legal sector can change the government's incentive to establish control. That sector must first be built. Investors pay to develop and install productive technologies, while verification providers maintain the evidence and financial commitments required for access to certified demand. The return to governing therefore depends on investment and on the institutions that sustain legal trade.

These decisions interact through the world market. When regions reform, prices change, altering the returns to research and the fiscal value of reform elsewhere. We characterize an equilibrium in which clean research and verification both receive resources and the expanding legal opportunity eventually induces intervention under a strict fiscal condition. The argument also identifies a countervailing force: removing illegal supply can increase the rents earned by those who remain. The dynamic analysis must accommodate this scarcity response as well as the possibility that a government waits or fails to establish control.

8.1 Maintaining credibility as production expands

Verification has both a current task and an inherited one. It must service new deliveries, resolve claims on earlier deliveries, and retain evidence of where old material originated. Let Q_t be newly serviced physical mass at date t , let claims remain open for T_H dates, and let D_t be the permanent record stock. Let $a_{\text{hist}} > 0$ weight service for unresolved historical cohorts. A conservative workload is

$$W_t = a_0 Q_t + a_{\text{hist}} \sum_{s=1}^{T_H} Q_{t-s} + a_D D_t, \quad D_{t+1} = D_t + n_t Q_t. \quad (97)$$

The coefficients are positive and the number of new records per physical unit is bounded. Current volume, unresolved cohorts and permanent evidence therefore contribute separately to the demand for verification services. Updating a claim changes its record without creating a new unit of mineral.

A funded research program employs $S_V > 0$ scientists to improve verification productivity Z_t . A release raises processing and storage efficiency by $g_V > 1$. Deployment, servers, replacement, maintenance and capital have positive costs. If physical throughput grows no faster than G^t , with $G > 1$, then

$$\frac{W_t}{Z_t} \text{ is uniformly bounded whenever } g_V \geq G.$$

With geometric growth, the accumulated record stock has the same order of growth as current throughput. Verification productivity that keeps pace can consequently service both new and old material with a bounded hardware fleet. Proposition F.3 establishes this bound while retaining the cost of research and deployment. It relies on productivity improvements applying to old-record storage: an irreducible expense per old record would prevent the same fixed-expenditure result.

Financial obligations expire according to their contracts. Each cohort has invoice-scaled liabilities backed by its own funded trust until the claims horizon closes. Procurement escrow pays the supplier, while separate assets secure sanctions and compensation. Tied collateral becomes available for new cohorts only after release. With at most T_H open cohorts and bounded invoice expenditure, outstanding principal is bounded. Initial infrastructure, inherited claims, new deposits, realized losses, reserve growth and the opportunity cost of capital all have explicit funding. Fees finance part of this system and an outside-good appropriation finances the balance. Proposition F.4 establishes its cash and resource feasibility. The evidence itself remains available after the financial obligation closes and after a region reforms.

8.2 Private returns and the direction of research

Profit opportunities influence which technologies investors develop [Acemoglu, 2002]. Their environmental consequences depend on whether research improves clean or harmful

production [Acemoglu et al., 2012]. Here the relevant opportunity is the investor’s own return under the actual selling and verification contracts. Ownership and the dates of payments therefore enter before the allocation of scientists is determined.

Scientific labor satisfies

$$S_V + \sum_i N_i s_{i,t} \leq S. \quad (98)$$

The mass N_i counts research owners of type i , separately from regions and auction buyers. An owner receives operating income from its own installation or an explicitly contracted licensing payment. Public access to a disclosed design gives adopters a production opportunity; it does not transfer their future tax base to the inventor.

A project commissioned at t operates its successful prototype at $t + 1$. The owner values success and failure at the actual world prices, subtracts the extra installation, service and collateral advances, and includes later payments and refunds attached to the affected cohort. In compact form,

$$J_{i,t} = \text{discounted own success cash flow} - \text{discounted own failure cash flow}. \quad (99)$$

Equation (192) records this dated difference in full. Disclosure can end a production lead after one interval while liabilities from the additional production remain outstanding. The original stock, menus, evidence and financial commitments stay in the success–failure comparison until their remaining consequences have been settled.

Some technologies make the operating increment particularly simple. At a common net quote, power-cost profit is linear in owned capacity. Replicating an independent complete retail installation likewise adds its optimized value, including the premium and retention decision appropriate to its mechanism. Production equipment, licensed access and institutions must all be financed. By contrast, a cost improvement within one nonlinear market changes the choice of output and retention and is valued by solving that new choice. Lemma F.5 states when later opportunities cancel, including the treatment of one-period prototype leases or recorded liquidation receipts. Outstanding cohort payments remain in the prize unless they too are identical after both histories.

After reserving verification scientists, a funded success payment can induce all remaining mineral research to serve a clean core. With owner mass N_* , each core owner then employs $s_* = (S - S_V)/N_*$. The sufficient condition compares the clean project’s lowest marginal private return at s_* , net of laboratory cost, with the highest marginal return of any competing project at zero effort and with zero. The comparison uses actual receipts and liabilities over the compact market domain. A finite funded success payment makes the inequality strict when required. The resulting standard grows at

$$g = 1 + \gamma_* (1 - e^{-\nu_* s_*}) > 1.$$

Theorem F.6 establishes this competitive allocation at endogenous prices. A finite verification productivity parameter allows verification to keep pace while $s_* > 0$. Wages, laboratories, success payments and installation all enter the outside-good budget. The

continuum construction uses the stated independent research outcomes and exact aggregate success fractions; a common scientific failure would instead be an aggregate shock.

Under this policy, mineral research remains concentrated on clean production throughout the transition. Government decisions still change project values and scientists' wages, but the strict return comparison keeps clean projects more attractive. Success payments require the funding specified in the budget.

8.3 Growth, market access and resources

The resources required by growth depend on the selling institution. In the common power-cost benchmark, bounded mineral expenditure can support rising physical output as productivity lowers prices. If clean capacity is of order g^t , throughput is at most of order $g^{t/(1+\eta)}$. Invoice-scaled liabilities keep required reserves bounded, while verification productivity can accommodate the greater physical workload. This is a conditional service and funding bound. A vulnerable origin still needs enough actual illegal stock for any false admission it supplies.

For the infinite common-market equilibrium, safe production and its funded certified expenditure expand together at scale $L_t = g^t$. Vulnerable capacity and a positive dirty bulk core stay fixed in absolute units, with a positive dirty expenditure floor. Sufficiently small positive vulnerable capacity can draw on available illegal stock at every date. Fraud and diligence remain endogenous there; the expanding safe flow reduces these origins' weight in aggregate contamination.

The nonlinear retail economy also uses clean portfolios of size $L_t = g^t$, with outside endowments and demand of the same order. Clean innovators own their portfolios. Public adopters install the disclosed standard with a one-date lag, paying for equipment and access. Fixed positive research cores are separate from the small, strictly positive population of reformable Selvas. Retail demand is connected to growing safe bulk supply and fixed vulnerable origins by the positive substitution term in the final consumer's utility.

Lemma F.2 establishes that all three normalized clearing equations, including the scale variable $1/L_t$, remain valid throughout the infinite path. Prices stay within the compact positive interval on which the local inventory and participation conditions hold. Growth in outside resources and demand finances the system and sustains access to these retail opportunities; optimal research generates growth in the clean standard. At fixed expenditure, indefinite replication could lower prices enough to change the retail inventory regime, so that path requires a different production analysis.

Outside households have the endowment required to buy final services, pay the lump-sum appropriations and retain positive numeraire consumption. Grants redistribute this endowment. Let $\delta_G, \delta_R \in (0, 1)$ be the one-period discount factors of governments and research owners, respectively. Both growing-income constructions impose $\delta_G g < 1$ and $\delta_R g < 1$, yielding uniform discounted tail bounds for normalized government and research values. In the fixed-expenditure benchmark the corresponding bounds use the ordinary

discount factors. These restrictions make future income and continuing institutional costs part of a feasible intertemporal allocation.

8.4 Government choices and their market consequences

The return to governing depends on the income a government can obtain through legal production. This connects the investment decision to the fiscal incentives studied by [Acemoglu \[2005\]](#) and to the complementarity between legal and fiscal capacity in [Besley and Persson \[2009\]](#). Governments also affect one another's returns. Where [Acemoglu et al. \[2015\]](#) study local state capacity through a municipal network, the interaction here operates through common mineral and final-product prices.

Let w_r measure regional scale. Within each normalized type, productive capacity, intervention cost, household measure and tax flows scale by w_r . Let μ_a^{reg} be the geographic measure of government type a . The ungoverned mass is $M_{a,t}^D = \int w_r 1_{\{r \text{ ungoverned}\}} d\mu_a^{\text{reg}}$. If fraction $u_{a,t}$ intervenes and success has probability P_a , the exact population transition is

$$\begin{aligned} M_{a,t+1}^D &= (1 - u_{a,t}P_a)M_{a,t}^D, \\ M_{a,t+1}^C &= M_{a,t}^C + u_{a,t}P_aM_{a,t}^D. \end{aligned} \tag{100}$$

Successful intervention changes production from the following date. The resulting supply determines new prices, certification composition, research prizes and government values. One atomistic government takes this public path as given. The actions of a positive population must collectively generate the path each government anticipates.

An ungoverned authority can accept the current protection payment and wait, or pay the intervention cost now. Failure leaves it ungoverned and able to try again. Success opens the legal tax base after installation, oversight and household support have been funded. For type a , let $k_a > 0$ be the normalized all-in intervention cost, $P_a \in (0, 1)$ its success probability, and $\hat{R}_{a,t}$ and $\hat{b}_{a,t}$ the normalized net legal fiscal flow and protection payment. These flows are measured per unit of regional scale and current income scale. Let $U_{a,t}$ and $V_{a,t}$ be the normalized values of a governed and an ungoverned region. The normalized government continuation factor is $\hat{\delta}_G = \delta_G$ at fixed expenditure and $\hat{\delta}_G = \delta_G g < 1$ in both growing-income classes. The values satisfy

$$\begin{aligned} U_{a,t} &= \hat{R}_{a,t} + \hat{\delta}_G U_{a,t+1}, \\ V_{a,t} &= \max\{\hat{b}_{a,t} + \hat{\delta}_G V_{a,t+1}, -k_a + \hat{\delta}_G[P_a U_{a,t+1} + (1 - P_a)V_{a,t+1}]\}. \end{aligned} \tag{101}$$

Waiting preserves the opportunity to intervene later. Reform changes future income, while historical claims remain the obligations of their original funded parties.

Theorem [F.7](#) establishes an infinite-horizon equilibrium of these government choices, research decisions and market prices. The proof constructs finite-horizon equilibria, permits mixing at genuine indifference between waiting and intervention, and uses compact normalized states and uniform continuation bounds to pass to the limit. Permanent records

and the finite queues of unresolved claims remain in the state. This consistency argument accounts for the price changes produced by collective reform; an individual government's discounted optimization alone would not determine those changes.

8.5 Reform and the rents of remaining illegal producers

A region's successful reform removes part of illegal supply. Holding clean supply and demand fixed, this raises the price available to the remaining illegal producers. The distinction between aggregate quantities and local rents is immediate in a simple equilibrium. With unit supply elasticity and dirty expenditure equal to one, $p_D = B_D^{-1/2}$. Reducing the supply coefficient from 1 to 0.81 lowers total dirty quantity from 1 to 0.9 but raises price to 10/9. A surviving unit-scale region expands production and sees profit rise from 0.5 to 50/81. Protection income can therefore increase in regions left outside reform even as aggregate illegal extraction declines.

The active nonlinear market has the same scarcity force. Its magnitude comes from final-output clearing and the cell's optimized profit. To determine whether clean expansion and changing demand outweigh it, compare supply and expenditure at the old price. In the scalar final-good market, price cannot rise when

$$p_t \sum_j (\hat{n}_{j,t+1} - \hat{n}_{j,t}) Z_j(p_t) \geq e_{t+1} - e_t.$$

The left side measures the value of the change in available output at the old price; the right side is the change in expenditure. Proposition F.8 establishes this comparison, and Corollary F.9 extends it to the fully coupled infinite path. A gradual decline in normalized retail spending toward a positive limit, large enough to dominate the maximum loss of dirty capacity and the coupled price effects, gives a falling retail price. Small positive dirty-core and Selva masses make these conditions compatible with the active mechanism. The result combines endogenous clean innovation with a funded demand policy.

The extraction consequence follows from the cell's physical input q_D , rather than its converted output Z_D . Muir physical output increases with price on the stated branch, while VCG physical output is constant on its strict inventory branch. The price condition therefore weakly lowers extraction by each unchanged dirty retail cell. The auxiliary bulk dirty price can move differently. A jurisdiction holding both retail and bulk activities must add their actual output and fiscal responses to assess its total change. The fixed dirty cores remain productive in this economy.

8.6 The decision to establish control

Eventually profitable intervention does not require monotone legal and illegal revenues. It requires a sufficiently attractive legal opportunity relative to the value of continued protection and the cost of acting. Suppose type a has a normalized funded net legal-

revenue floor $\underline{R}_a > 0$ and normalized protection payments tend uniformly to zero over feasible population paths. The strict condition

$$k_a(1 - \hat{\delta}_G) < \hat{\delta}_G P_a \underline{R}_a \quad (102)$$

makes intervention optimal after a finite date for every region that remains ungoverned. Theorem F.10 establishes this comparison using the government's full value of waiting, including its ability to act later and to try again after failure.

The growing legal standard supplies the economic force behind the condition. In both growth constructions, a remaining illegal installation has bounded absolute income, whereas a legal adopter installs the growing public standard. Funded legal fiscal income is a positive fraction of L_t , and protection income divided by L_t tends to zero. Positive but sufficiently small intervention and support costs satisfy the inequality on an open parameter set. The additional retail demand policy can also reduce absolute dirty retail rents, but expansion of legal opportunity already supports the normalized fiscal comparison.

Waiting and persistent nonintervention remain possible outside these conditions. A government can optimally postpone action while protection income is high, and an unsuccessful attempt can be followed by another. If protection remains at least as attractive as the legal public flow and intervention costs are positive, waiting forever can be optimal. With zero probability of success, intervention yields no gain. In a fixed-expenditure market, the disappearance of actual protection income must be established separately; a fixed dirty expenditure share does not provide it. These cases show why the fiscal opportunity, success probability and option to wait all enter the intervention criterion.

8.7 Household livelihoods and protected sites

Establishing control changes the production opportunities available to households. Each market allocation is returned to them through the regional implementation correspondence. Legal tax revenue first pays for oversight and the announced livelihood contract. Let f_S and f_L bound local-service and livelihood-support costs per unit of scale. If legal operating surplus per unit of scale is at least $\underline{\pi}$, then


$$f_S + f_L < \tau \underline{\pi}$$

finances local services f_S , livelihood support f_L and a positive remaining fiscal return. Support is committed from the implemented allocation and does not change a household's marginal return to its subsequent output choice. Advance reserves finance the interval before taxes arrive. Preparation includes access to legal production and, where required, the licensed retail technology. The legal opportunity is thus available to households as a financed allocation.

A conservation result further requires that legal production use a safe physical technology. Here legal installations occupy specified noncritical sites and cause zero extraction disturbance in the protected forest; fixed research cores also lie outside that habitat. Once intervention becomes strictly optimal, the remaining ungoverned Selva mass declines geometrically. Aggregate disturbance in the protected sites tends to zero, and every positive tolerance is reached after a finite date. Imperfect success in an atomless population allows some ungoverned mass at every finite date, even though each individual region eventually succeeds almost surely.

For irreversible damage, the time taken to establish control matters. Proposition [F.11](#) bounds disturbance both before intervention begins and during repeated attempts. A feasible conservation policy must fit this transition within its cumulative-disturbance budget. Reform leaves old mineral's origin intact, and later legal production cannot restore physical losses treated as irreversible. The economic route to conserving the protected forest therefore combines incentives to establish control with funded household livelihoods, a safe production location and action early enough to respect that physical budget.

A. Scarcity can reward the illegal producers that remain



Dirty expenditure and clean supply are fixed.

B. A funded route from legal opportunity to protected-site conservation





Figure 6: Scarcity and the conditional route to conservation. Panel A holds dirty expenditure and clean supply fixed: less aggregate illegal extraction can coexist with higher surviving-region rents. In panel B, navy dash-dot arrows show the conditional path and repeated attempts; the gold frame highlights the fiscal comparison. Intervention becomes optimal after a finite date under the stated conditions, while protected-site disturbance declines through repeated attempts. Some ungoverned mass can remain at every finite date, and fixed dirty cores remain productive. Historical origin and liabilities persist; irreversible losses remain subject to the cumulative-disturbance budget.

9 Conclusion



Market access, investment and funded intervention connect exterior policy to livelihoods and safe production sites, with protected-site disturbance declining gradually.

Prices coordinate exchange, but the institutions governing that exchange determine how a restriction reaches production. When an intermediary controls both mineral exports and essential imports, a change in the purchases of distant consumers operates through its remaining outlets and the household choices it can induce. We have established conditions under which these responses preserve extraction even as regulated transactions and intermediary receipts change.

The whole-buyer VCG calculation identifies which regulated orders affect payments and the return to additional stock. Retention and unaffected buyers can sustain that return. The seller-payoff bound specifies the efficient-outcome class in which the mechanism is optimal; the comparison across inventories determines whether a restriction reduces extraction. In both settings, the relevant comparison includes every feasible production choice and the household allocation required to implement it.

Credible provenance changes the supplier's opportunities by allowing buyers to condition demand on origin. Credibility requires resources: inspectors must be paid, evidence

preserved and liabilities enforced. The people who can facilitate fraud must also have reason to refuse it or report it. These obligations continue after the initial sale. When they are funded, certified demand can reward legal production, although illegal material may continue to find buyers elsewhere.

Political incentives create a further obstacle. Reform can reduce aggregate illegal supply while raising the price and protection rents of suppliers that remain. Governments therefore need a sufficiently valuable legal alternative to prefer intervention. In the equilibrium we construct, funded clean innovation expands that alternative. Legal revenues eventually outweigh protection income and intervention costs under the stated fiscal condition, even after accounting for the option to wait and the possibility of failure.

The environmental result depends on what successful intervention makes possible. Legal production takes place at safe sites, household livelihoods are funded and the remaining ungoverned population declines. Disturbance at the designated protected sites converges to zero. Where damage is irreversible, the transition must also satisfy the cumulative physical constraint. The result thus requires both an incentive to establish control and resources to sustain the resulting allocation.

These findings identify the evidence needed to assess a sourcing policy: the outlets available to displaced mineral, the local terms governing household work, the cost of verifying origin and the returns to legal and illegal activity that governments can obtain. A change in purchases alone does not reveal these responses. Conservation becomes feasible when policy changes the economic opportunities of the actors who produce, trade and govern the mineral.

Notation

The paper introduces each object where it is first needed. These tables collect the notation in one place. The first table collects recurring notation. The following tables list, section by section, the symbols specific to each part of the argument. A dagger ([†]) marks a symbol that also carries another meaning elsewhere; the last table lists these cases.

Global notation

Symbol	Meaning	Sections
q	Physical mineral stock delivered by a region	2–7
$\bar{q}$	Size of a standardized package	2–3
θ_B	Procurement price posted by the gate for mineral	2–3
$W(q)$	Gross value to the gate of delivering stock q downstream; W_r for region r	2–7
$\mathcal{C}(q), \Gamma(q)$	Least cost of implementing output q ; set of cost-minimizing implementations	2–5
Π	Profit of the gate	2–3, 5–7
τ	Compliance cost per standardized unit borne by regulated buyers [†]	2–3
p_D	Price in the outside, unverified market	4, 6–8
h	Certified quote, the commodity invoice in the certified market [†]	4, 6–8
x	Price ratio p_D/h between the outside and certified markets [†]	6–7
ι	Illegal share of certified mineral	4, 6, D
Ω	Public market state, including committed continuation arguments	4–5, 7
π	Share of fraudulent admissions; posterior probability of fraud [†]	4, 6
η	Elasticity of regional output with respect to its unit receipt [†]	5–8
E_C, E_D, E_R	Expenditure on certified, outside and retail purchases	4, 6–7
X, Y	Certified and outside purchases	6–7
L_t	Scale of safe supply and retail installations at date t [†]	7–8
k	Origin of mineral [†]	4, 6
t	Date	4, 7–8
Subscripts C, D	Certified (clean) and outside (dirty)	throughout

Section 2: Selva, production behind an armed chokepoint

Symbol	Meaning
$\mathcal{J}, v$	Villages and production sites of Selva; a village
g	The chokepoint, or gate, through which all mineral leaves the region [†]
$\mathcal{G} = (\mathcal{V}, \mathcal{L})$	Directed network of usable trade routes
$\mathcal{E}$	Designated set of ecologically critical sites (Assumption 2.4)
$\mathcal{A}_r^{\text{safe}}$	Safe physical allocation set of site r
$\mathcal{V}_J, A$	Potential villages; active network chosen by the gate
$u_v, \bar{U}_v$	Household utility; outside utility
C_v, c_v	Consumption basket; subsistence level
r_v, b_v, b_{0v}	Leisure; retained cash; retained-cash utility allowance
η_v, β_v	Utility weights on leisure and retained cash [†]
α	Share of imported necessities in the basket
a_v	Opening wealth of the household [†]
l_v^x, l_v^m	Labor in local-good production and in mining
F_v, G_v	Local-good and mineral technologies
g_v	Purchased mining input
K_v, A_D	Inherited equipment; dirty mining technology
$p_v, \bar{p}_v$	Local-good price; consumption-basket price ($\bar{p}_{hr}$ in Appendix C)
p_v^z, p_v^g	Prices of imported necessities and mining inputs
c_v^z, c_v^g	Delivered resource costs of necessities and inputs
μ_z, μ_g	Margins of the gate on necessities and inputs
w_v	Wage
t_{vw}, δ_{vw}	Shipment from v to w ; transport loss factor
$e_v, \Psi_v(e_v)$	Coercive effort; its real cost [†]
$\Delta_v(e_v)$	Deterioration of the outside option under coercion [†]
w, d, B	Distress benchmark: mining earning power; time to produce one basket's local complement; price of the imported component [†]
q, Z, I	Regional mineral, imported necessities and mining inputs
$z, \mathcal{Z}$	An implementation (terms, active set, coercion, equilibrium); the set of supportable equilibria
$T(z)$	Transport and coverage bill [†]
$n, \mathcal{N}(q)$	Offered capacity; allowable capacities at output q

Section 3: Selective regulation and the return to extraction

Symbol	Meaning
θ_E	Price at which laundered mineral sells outside the region
$\mathcal{R}, \mathcal{U}$	Regulated and unaffected buyers [†]
$\mathcal{I}, N$	Buyers; their number
$H(y)$	Retention value of unsold stock y , through laundering, storage or disposal [†]
$\bar{k}_i, \omega_i, \nu$	Capacity and private type of buyer i ; type law
J_i^0	Operating value of buyer i before payment for the auctioned mineral
b_i^k	Marginal value of buyer i 's k th package
$v_i(k; \omega, \tau)$	Value of k packages to buyer i , net of compliance cost
Q, r_q	Number of whole packages in stock q ; remainder
ℓ, s, d	Laundering, storage and disposal of retained stock [†]
$V(\ell), S(s)$	Real cost of laundering; value of storage
S_n, p_i	Surplus with n offered packages; payment of buyer i
a	Stock outside offered capacity, $a = q - n\bar{q}$, including the fractional residual [†]
$\rho_j(y)$	Marginal retention value of package j above stock y ; seller orders are $\rho_1(a), \dots, \rho_n(a)$
$\rho(q)$	Smallest seller order: $H(q) - H(q - \bar{q}) = \rho_n(a) = \rho_1(q - \bar{q})$ for $n \geq 1$
$D_i(y)$	Retention value of the winning buyer's block, equal to $\sum_{u=1}^{k_i^*} \rho_u(y)$
N_i	Newly filled rival-order value, equal to p_i , only in the proof of Lemma 3.5 [†]
p^R	Index of regulated buyers' unit values, $\theta_E - \tau$
L, m	Actual laundering; packages sold
$e_{\mathcal{U}}$	Depth of unaffected rival demand
$\Xi = (z, n), \mathcal{X}$	A full implementation; its feasible set
$\bar{\chi}$	Expected exposure of payments to regulation
$\mathcal{H}, \Lambda, \gamma$	Hessian of profit; transmission terms in the local result

Section 4: Provenance, private incentives, and the cost of credibility

Symbol	Meaning
$\mathcal{M}, m(\cdot)$	Physically identified mineral portions; their mass
$c, v(c), \psi(c)$	A live claim; its mass; its physical backing
S_t	Live claims after settlement event t

Symbol	Meaning
m_{loss}	Recoverable mineral content lost during a transfer
$q_{D,k}, d_k$	Illegal stock at origin k ; attempted diversion [†]
$\mathcal{O}_t^{\text{rev}}$	Origins whose certification has been revoked [†]
$\mathcal{O}_t^{\text{susp}}$	Origins under temporary suspension [†]
A_t, X_t	Cumulative admitted mass; cumulative exits
V_k, U_k	Live certified mass from origin k ; its illegal component
T_H	Discovery and claims horizon
e_j, c_j	Investigation effort of purchaser j ; its cost
$c_j^{\text{sil}}, c^{\text{sil}}$	Channel-specific and homogeneous silence costs; $c_{\min}^{\text{sil}}$ is their positive lower bound
j_{new}	Newly added discovery channel, outside the existing set J
r_0, r_j	Baseline discovery probability; discovery through channel j
P_{-j}, P_n	Discovery probability through other channels; with n channels [†]
$L_{E,j}, L_{L,j}$	Purchaser's loss after an early report; after later discovery
$G_{\Omega}(a), K_J^F(a)$	Gain from a fraudulent action a ; funded cost of admission and suppression
$\iota^{\text{in}}, \iota^{\text{live}}$	Illegal share at admission; in the live certified stock
R_U, R_A	Illegal and admissible mass later removed
ϕ, E_0, F	Monetary fee per admitted unit; outside funds; recurring expenditure [†]
φ	Fraction of certified expenditure collected as fee revenue [†]
B_0, D	Organizer's bond; total bonds
κ	Cost parameter of buyer effort

Section 5: Regional production and world supply

Symbol	Meaning
$(\mathcal{R}_g, \mu_g^{\text{reg}})$	Regional population of group g ; its mass is m_g . Appendix C suppresses g in μ^{reg}
$r_g(\Omega)$	Unit receipt of group g
$\Gamma_r, \mathcal{C}_r$	Implementation correspondence and least cost of region r
Q_g, B_g	Group output; its scale coefficient
β_r	Supply scale of region r

Section 6: Certification and world-market equilibrium

Symbol	Meaning
B_{Ck}, B_{Dk}	Clean and dirty supply coefficients of origin k
Q_C, Q_D	Clean and dirty output, $Q_{Ck} = B_{Ck}h^\eta$, $Q_{Dk} = B_{Dk}(xh)^\eta$
c_k	Upper bound of the forgery-cost draw
r_k	Probability of unsuppressible late discovery
b_k, hb_k	Invoice-normalized collectible bond coefficient; total collectible bonds
hl_{Ek}, hl_{Lk}	Buyer losses after an early report; after late discovery
π_k, e_k	Fraud share and diligence at origin k
X_k, w_k	Admissions from origin k ; their share
$F_{cw}(x)$	Admission coefficient: $X = h^\eta F_{cw}(x)$
$H_{cw}(x)$	Diversioin coefficient: $Y = h^\eta [B_D x^\eta - H_{cw}(x)]$
ϕ	Invoice-relative service-fee coefficient; the monetary fee per unit is $h\phi^\dagger$
$T(x)$	Certified-expenditure coefficient: $E_C = h^{1+\eta}T(x)$
$J(x)$	Ratio of outside to certified expenditure implied by the common-price clearing equations
$\underline{x}, \bar{x}$	Bounds of the studied price-ratio branch

Section 7: Retail market power and the world price

Symbol	Meaning
n_r	Mass of production and retail cells in region r
$n_{j,t}, \hat{n}_{j,t}$	Absolute and normalized retail cell masses: $\hat{n}_{j,t} = n_{j,t}/L_t$; $\hat{n}_t$ is the vector
$Z_R(p), E_R, p$	Final retail supply; retail expenditure; retail price
$q_j(p), \pi_j(p)$	Output and profit of a cell of type $j^\dagger$
k	Cost parameter of a cell [†]
$K(x)$	Generalized certified-contract cost divided by its commodity invoice h
F_{cw}, H_{cw}	The admission and diversion coefficients defined in Section 6
λ_t	Inverse scale, $1/L_t$

Section 8: Technical change and the incentive to govern

Symbol	Meaning
Q_t, D_t	Newly serviced physical mass; permanent record stock
T_H, a_{hist}	Contractual claims horizon; workload per unit in an unresolved historical cohort
S_V, Z_t, g_V	Verification scientists; verification productivity; efficiency gain per release
w_r	Scale of region r
μ_a^{reg}	Geographic measure for government type a
$\hat{B}_t^C, \hat{B}_t^D$	Capacities divided by L_t ; equal to normalized installation masses only in the retail interpretation
$\hat{n}_{C,t}, \hat{n}_{D,t}$	Normalized clean and dirty retail installation masses
$M_{a,t}^D, M_{a,t}^C$	Ungoverned and controlled mass of type a
$u_{a,t}$	Fraction of ungoverned governments that intervene
k_a, P_a	Normalized all-in intervention cost; success probability for government type a
$\hat{\delta}_G$	Normalized government continuation factor: δ_G at fixed expenditure and $\delta_G g$ with growing income
$\underline{R}_a$	Positive floor of normalized net legal fiscal revenue for type a
δ_G, δ_R	Primitive discount factors of governments and research owners
$\hat{R}_{a,t}, \hat{b}_{a,t}$	Normalized net legal fiscal flow and protection payment
$U_{a,t}, V_{a,t}$	Normalized government values when controlled and ungoverned
$\underline{\pi}$	Lower bound on legal operating surplus per unit of scale [†]
f_S, f_L	Local-service and livelihood-support cost bounds; $f_{S,a}, f_{L,a}$ for type a

Symbols with several meanings

Symbol	Meanings (section)
ℓ	Laundering quantity (3); research-payment bound ℓ_{iu} (Appendix F)
d	Time per basket (2); disposal (3); attempted diversion d_k (4)
P	Discovery probabilities (4); probability of successful intervention P_a (8)
R	Benchmark constant (2); regulated buyers $\mathcal{R}$ (3); removed mass (4); regional space $\mathcal{R}_g$ (5); legal fiscal revenue (8)
h	Certified quote (4, 6–8); research success function $h_i(s_i)$ (8, Appendix F); integrable bound $h(r)$ (Appendix C); retained-conversion productivity (Appendix E)

Symbol	Meanings (section)
x	Local good and endogenous vector (2); price ratio (6–7)
H	Retention value (3); added recurring service cost (4); household income H_r (5); diversion coefficient H_{cw} (6–7)
τ	Compliance cost (2–3); reservation of a reporter τ_j (4); tax rate on legal surplus (8)
π	Fraud share (4, 6); cell profit π_j (7); legal surplus per unit of scale $\underline{\pi}$ (8)
η	Utility weight on leisure η_v (2); output elasticity (5–7)
k	Package index (3); origin (4, 6); cell cost (7); intervention cost (8)
a	Household wealth (2); stock not offered (3); fraudulent action (4); government type (8)
L	Laundering (3); purchaser losses L_E, L_L (4); scale L_t (7–8)
$\mathcal{O}$	Suspended or revoked origin sets (4); owner operating cash receipts $\mathcal{O}_{i,t+1}^{s/f}$ (Appendix F)
Ψ	Coercion cost $\Psi_v(e_v)$ (2); cell receipt term in $p\Psi - C$ (Appendix F)
Δ	Outside-option deterioration $\Delta_v(e_v)$ (2); profit gap Δ_0 (3); finite or success-minus-failure differences (Appendix F)
ϕ	Monetary unit fee (4); invoice-relative service-fee coefficient (6–8)
φ	Expenditure fee share (4); auxiliary labor-response function φ_i (Appendix A)
N_i	Newly filled rival-order value in one proof (3); research-owner mass (8, Appendix F)
g	Chokepoint and mining input g_v (2); group of regions (5); growth rates g, g_V (8)
T	Transport bill (2); silence set (4); fraud gain T_k and certified-expenditure coefficient $T(x)$ (6)

A Regional implementation: full derivations

A.1 Economic bounds for the smooth linked class

The result concerns a nonempty equilibrium graph under Assumption 2.9. It supplies compactness on positive-output cost sublevels and attainment of the cost fiber; it does not infer existence of every heterogeneous equilibrium from individual concavity. The argument keeps the full owner income, both import margins, costly coercion and route complementarity. Here $a_x = 1 - \alpha$ and N_m denotes the number of mining-capable villages.

Complete proof of Theorem 2.10. Let $H = M + \sum_v a_v$. The exact resource account implies $\alpha c_v^z C_v, c_v^g g_v, b_v \leq H$ individually. Thus C_v, g_v, b_v have explicit upper bounds, and $r_v \leq 1$. Put $\underline{U}_v = \bar{U}_v - \Delta_v(\bar{e}_v)$. Participation gives

$$C_v - c_v \geq s_v := \frac{e^{\underline{U}_v}}{(H + b_{0v})^{\beta_v}} > 0, \quad r_v \geq t_v := \left[\frac{e^{\underline{U}_v}}{(H/(\alpha c_v^z) - c_v)(H + b_{0v})^{\beta_v}} \right]^{1/\eta_v} > 0$$

If a displayed upper bound cannot accommodate these lower bounds the set is empty, contrary to the premise. From cash KKT and consumption/leisure KKT,

$$\underline{w}_v := \eta_v \alpha c_v^z s_v \leq w_v = \eta_v / (\lambda_v r_v) \leq \frac{\eta_v (H + b_{0v})}{\beta_v t_v} =: \bar{w}_v$$

For a capable miner both factors are positive: starting both at a small ϵ yields a benefit of order $\epsilon^{\gamma_v + \kappa_v}$ against linear input and time costs; the former dominates at zero. The interior factor conditions therefore give

$$\theta_B G_v = w_v l_v^m / \gamma_v, \quad p_v^g g_v = \kappa_v \theta_B G_v \quad (103)$$

Hence $\theta_B q \leq D := \sum_{A_v^m > 0} \bar{w}_v / \gamma_v$. The external cash account and (17) imply $\theta_B q \geq d_A$. The resource input bounds and $l_v^m \leq 1$ give $q \leq \bar{q} := \sum_v A_v^m (H / c_v^g)^{\kappa_v}$. Consequently

$$0 < d_A / \bar{q} \leq \theta_B \leq D / q_0 < \infty$$

For nonnegative profit, $C(z) \geq d_A$ and $W \leq \bar{v}q$ imply the claimed automatic output floor. There is also an economic upper cost bound without assuming bounded output first. Let $G_T = \sum_v g_v$, $c_g^{\min} = \min_v c_v^g$, $a_T = \sum_v a_v$, $A_T = \sum_v A_v^m$ and $\bar{\kappa} = \max_v \kappa_v < 1$. For $G_T \geq 1$, $q \leq A_T G_T^{\bar{\kappa}}$ and $C(z) \geq c_g^{\min} G_T - a_T$. If additionally $G_T \geq 2a_T / c_g^{\min}$, nonnegative profit implies $(c_g^{\min} / 2) G_T \leq \bar{v} A_T G_T^{\bar{\kappa}}$. Thus all such choices have

$$G_T \leq \max \left\{ 1, \frac{2a_T}{c_g^{\min}}, \left(\frac{2\bar{v}A_T}{c_g^{\min}} \right)^{1/(1-\bar{\kappa})} \right\}$$

This bounds q , hence $C(z) \leq W(q) \leq \bar{v}q$, and supplies one finite sublevel for every choice able to beat exit.

At least one of the N_m capable miners has $G_v \geq q_0/N_m$, hence $g_v \geq [q_0/(N_m A_v^m)]^{1/\kappa_v}$. Equation (103) bounds its input quote above. Taking the largest of these finitely many bounds bounds the common margin μ_g and thus every input quote above. To bound them below, eliminate mining labor from the two factor conditions:

$$(p_v^g)^{1-\gamma_v} = \kappa_v \theta_B A_v^m \left(\frac{\gamma_v}{\kappa_v w_v} \right)^{\gamma_v} g_v^{\gamma_v + \kappa_v - 1}$$

Use the positive lower bound on θ_B , the upper bound on w_v , and the upper bound on g_v . The final exponent is negative, so the right side has a strictly positive lower bound. Apply it to a mining-capable village with minimal c_v^g ; every other posted input price is at least that price. This is the point at which strict decreasing returns removes the earlier fixed-proportion offset ray.

The cash account also yields $\alpha \mu_z \sum_v C_v \leq D + \sum_v a_v$, bounding the nonnegative necessity margin. Consumption KKT bounds each local price above, since $\lambda_v \geq \beta_v/(H + b_{0v})$, $C_v - c_v \geq s_v$ and $\bar{p}_v = a_x p_v + \alpha p_v^z$. Physical clearing implies $\sum_v F_v \geq a_x \sum_v c_v$. Some producer therefore supplies at least $a_x \sum_v c_v/N$. Its labor is bounded below by the inverse of its strictly increasing production function. The local production FOC $p_v F'_v = w_v$ then bounds that producer's price strictly above zero. For every other node, use a directed path from that node to this producer: trade zero-profit inequalities imply that the producer's price is no larger than the other node's price times the finite path iceberg product. Thus every local price is bounded away from zero. Finitely many producers and paths give a uniform bound. Positive-price zero-profit trade cannot have a cycle with a strict iceberg loss. Costless cycles can be removed while preserving every village balance and trade complementarity. Each allocation therefore admits an acyclic shipment witness, bounded by production and finitely many paths. The set of acyclic supports is a finite union of closed sets; its witness graph is closed and bounded. A minimum-flow convention can subsequently choose a representative. Its continuity is unnecessary because the theorem projects out shipment choices, and costs here depend on the preserved physical accounts.

All effective quantities and quotes now lie in a compact box separated from the survival and relevant price singularities. Positive-capability factor choices are also separated from zero by their factor equations with prices and wages in these bounds. Continuous KKT, market-clearing, participation and cost conditions have a closed graph, including cash and trade corners. Concavity makes those conditions sufficient for each household and firm. Limits are competitive equilibria. A minimizing sequence at fixed positive output lies in a sublevel furnished by any feasible candidate, so it attains its infimum. Intersecting with a closed predetermined selection graph preserves the argument; intersecting with an arbitrary discontinuous choice need not. $\square$

A.2 A unique linked price with trade-dependent survival

Village 1 produces only $A_x l_1^{\nu_1}$ local goods. Village 2 has $A_m l_2^{\gamma} g^{\kappa}$ mining, with $\gamma + \kappa < 1$. There is zero opening wealth and costless two-way local trade. Imported necessity basket

price is common $B = \alpha p^z > 0$. Let $a_x = 1 - \alpha$, $P = a_x p + B$, and eliminate the purchased input:

$$g = a_g l_2^{\nu_2}, \quad a_g = (\theta_B A_m \kappa / p^g)^{1/(1-\kappa)}, \quad \nu_2 = \gamma / (1 - \kappa), \quad D = (1 - \kappa) \theta_B A_m a_g^\kappa$$

The owner's net mining income is $D l_2^{\nu_2}$, including its fixed-factor profit. Household i has positive subsistence c_i , leisure weight η_i and valuable cash as before. Define

$$p_- := \frac{B c_1}{A_x - a_x c_1}, \quad p_+ := \frac{D / c_2 - B}{a_x}, \quad t_1(p) = \frac{p A_x}{a_x p + B}, \quad t_2(p) = \frac{D}{a_x p + B}$$

Require $A_x > a_x c_1$, $0 < p_- < p_+$, and the primitive cash inequalities

$$b_{01} / \beta_1 > A_x p_+, \quad b_{02} / \beta_2 > D \quad (104)$$

Theorem A.1 (Unique linked price from a monotone survival equation). *On (p_-, p_+) let $l_i(p) \in (\nu_i / (\nu_i + \eta_i), 1)$ solve*

$$t_i(p) l_i^{\nu_i-1} [(\nu_i + \eta_i) l_i - \nu_i] = \eta_i c_i \quad (105)$$

Define $f(p) = l_1(p)^{\nu_1} / l_2(p)^{\nu_2}$ and $k = a_x D / (A_x B)$. If $f(p_+) < k < f(p_-)$, interpreting endpoint limits, there is exactly one strictly surviving competitive equilibrium. Its price is the unique root $f(p) = k$, its cash is zero, and its household/firm choices are the stated optimal choices. The selection is independent of downstream policy and continuous on the strict primitive domain. Both singleton coverage sets fail survival. Participation must additionally be checked against the given outside utilities/coercion bounds.

Proof. Any surviving equilibrium needs $t_i(p) > c_i$, which gives exactly the displayed price interval. Income is bounded above by $A_x p_+$ for household 1 and D for household 2. At positive cash, KKT would give $\lambda_i \leq \beta_i / b_{0i}$, whereas $1 = \lambda_i P(C_i - c_i) < \lambda_i \text{income}_i < 1$ by (104). Thus cash is zero in every surviving equilibrium; this is proved from primitives, not imposed as an omitted choice.

Substitute $C_i = t_i l_i^{\nu_i}$ in $\log(C_i - c_i) + \eta_i \log(1 - l_i)$ and differentiate. Its strictly concave owner problem yields (105). The function $\varphi_i(l) = l^{\nu_i-1} [(\nu_i + \eta_i) l - \nu_i]$ increases strictly for $l > \nu_i / (\nu_i + \eta_i)$, because

$$\varphi_i'(l) = \nu_i l^{\nu_i-2} [(\nu_i + \eta_i) l + 1 - \nu_i] > 0$$

It ranges from zero to η_i as l increases to one. Since $t_i > c_i$, the root is unique and $dl_i/dt_i = -\varphi_i / (t_i \varphi_i') < 0$. The real earning power t_1 increases strictly with p , whereas t_2 decreases strictly. Therefore f decreases strictly and continuously. The endpoint inequalities give exactly one root by the intermediate value theorem. Strict monotonicity and continuous primitives give continuity of the root; differentiability holds where the displayed derivatives are used.

Local goods clearing is $A_x l_1^{\nu_1} = a_x(C_1 + C_2)$. Substitute the two household budgets and multiply by P to reduce it to $A_x B l_1^{\nu_1} = a_x D l_2^{\nu_2}$, exactly $f = k$. All income, labor and input conditions are household/firm optima; shipments clear with zero trade profit. Without village 2, the summed budget cannot finance imported necessities; without village 1 there is no local complement. Neither singleton survives. The argument does not select a cartel-favorable equilibrium from a multiplicity: the competitive price is unique. $\square$

A.3 Effective terms and nominal multiplicity

Consider a village with fixed-proportion mining $G(l, g) = A \min\{l, g\}$ and local production $F(l^x) = (1 - \alpha)l^x/d$, where $d, A > 0$ and $dc < 1$. There is no opening wealth. Positive input prices imply $g = l = q/A$. Define net mining earning power $w = A\theta_B - p^g$ and $B = \alpha p^z \geq k = \alpha c^z > 0$. The exact owner problem is

$$\begin{aligned} \max \quad & \log(C - c) + \eta \log r + \beta \log(b + b_0) \\ \text{subject to} \quad & dC + l + r = 1, \quad BC + b = wl, \\ & C > c, \quad r > 0, \quad b, l \geq 0 \end{aligned} \tag{106}$$

Survival requires $w > 0$. Eliminating l gives the full-income constraint $(B + wd)C + wr + b = w$. With $D = 1 + \eta + \beta$ and $M = w + b_0 - (B + wd)c$, the positive-cash solution is

$$C - c = \frac{M}{D(B + wd)}, \quad r = \frac{\eta M}{Dw}, \quad b + b_0 = \frac{\beta M}{D} \tag{107}$$

It applies when $b > 0$ and $l > 0$. At $b = 0$, the two-good problem uses resource $w - (B + wd)c$ and weights $1, \eta$; its solution is accepted exactly when $\lambda \geq \beta/b_0$. With no wealth, a zero-mining allocation cannot pay for the positive imported basket component.

Proposition A.2 (Attained effective terms, nonunique gross quotes). *With bounded coercion, finite outside utility and continuous nonnegative transport and coercion costs, every nonempty cost sublevel in (106) is compact in physical allocations and effective terms (w, B) . Every nonempty output fiber attains its cost minimum. Its nominal implementation set contains*

$$(\theta_B, p^g) \longmapsto (\theta_B + t, p^g + At), \quad t \geq 0$$

whenever those terms are allowed. This invariance can hold with strictly positive retained cash.

Proof. The consolidated resource cost is $kC + c^g l + b + T + K$. On a sublevel bounded by M_0 , cash is bounded by M_0 , consumption by $1/d$, and leisure by one. Participation supplies $U_{\min} = \bar{U} - h(\bar{e}) > -\infty$. Consequently

$$C - c \geq \frac{e^{U_{\min}}}{(M_0 + b_0)^\beta} =: s_* > 0, \quad r \geq \left[\frac{e^{U_{\min}}}{(1/d - c)(M_0 + b_0)^\beta} \right]^{1/\eta} =: r_* > 0$$

Consumption and leisure optimality imply $B/w = r/[\eta(C - c)] - d > 0$, and cash complementarity gives $w \leq \eta(M_0 + b_0)/(\beta r_*) =: w^*$. The necessity floor then gives $w \geq k\eta s_* > 0$ and $B \leq w^*/(\eta s_*)$. The budget also bounds mining labor away from zero: $l \geq k(c + s_*)/w^*$. All these bounds are economic consequences of resources and utility.

The linear resource constraints, household optimality, participation and cost conditions have closed graphs inside this region. A convergent minimizing sequence therefore attains the minimum. A finite nominal representative is $p^g = c^g$, $\theta_B = (w + c^g)/A$ and $p^z = B/\alpha$. Under the displayed price transformation, w and B are unchanged. The increased mineral payment tq exactly equals the increased input payment $Atg = tq$. All household choices and consolidated costs are unchanged. The positive-cash formula permits $\beta M/D > b_0$, so the invariance is compatible with valuable, strictly positive cash. A nominal normalization selects a representative of these terms; it does not prove their economic uniqueness. $\square$

A.4 An implemented homogeneous output-cost branch

The global price analysis uses a transparent specialization in which the regional cost envelope has a homogeneous positive-output branch. This subsection supplies a full household-owner implementation of that branch. It is a restricted institution within the regional framework: necessities are sold at pass-through terms, productive-input subsidies are excluded, and mining uses inherited machinery with variable purchased inputs and no variable household labor. The labor-using jungle economy remains the general regional model.

A nonmining sector or endowed outside-numeraire asset supplies household wealth sufficient for a strictly surviving consumption plan. Its consumption opportunities and prices are fixed in this specialization; the local-good component can be supplied by a competitive background sector at its given resource price. These endowments are owned resources, not unfunded cartel transfers. Indirect household utility is strictly increasing in additional operating income. Consequently, at given mineral and input quotes, its choice of productive input maximizes the net income of its owned mining firm. Current consumption and saving divide that income without changing the firm's input choice.

Let installed mining productivity be $a(A) > 0$ and technology be

$$q = a(A)g^{\alpha_m}, \quad 0 < \alpha_m < 1 \quad (108)$$

The gate posts $\theta_B > 0$ and $p_g \geq c_g > 0$ and sets the necessity margin to zero. Its other regional costs on this branch are zero; an explicit fixed operating requirement could be added as a separate branch cost but is not needed for the result. The owner chooses $g \geq 0$ to maximize $\theta_B a(A)g^{\alpha_m} - p_g g$. The gate's net procurement cost is the gross mineral payment less its productive-input margin.

Proposition A.3 (An attained power cost envelope). *Every target $q > 0$ has an implemented cost-minimizing input quote $p_g = c_g$ and a positive procurement quote. Its envelope and implementing quote are*

$$\begin{aligned}\mathcal{C}(q; A) &= \frac{c_g}{\alpha_m} a(A)^{-1/\alpha_m} q^{1/\alpha_m}, \\ \theta_B(q; A) &= \frac{c_g}{\alpha_m} a(A)^{-1/\alpha_m} q^{1/\alpha_m - 1}\end{aligned}\tag{109}$$

Household operating rent is $(1 - \alpha_m)c_g g / \alpha_m$ and is included in this implementation cost. Set $\alpha_m = \eta / (\eta + 1)$ for $\eta > 0$ and choose

$$a(A) = \left(\frac{c_g}{\kappa \alpha_m^2} \right)^{\alpha_m} A^{1 - \alpha_m}, \quad \kappa > 0\tag{110}$$

Then the envelope becomes

$$\mathcal{C}(q; A) = \frac{\kappa}{1 + 1/\eta} A^{-1/\eta} q^{1 + 1/\eta}\tag{111}$$

Proof. Strict concavity of net firm income gives a unique positive input:

$$\theta_B a(A) \alpha_m g^{\alpha_m - 1} = p_g, \quad \theta_B q = p_g g / \alpha_m$$

To implement a target q , technology fixes $g = (q/a(A))^{1/\alpha_m}$ and this condition fixes the necessary positive procurement quote at every allowed p_g . The gate's net bill is

$$\theta_B q - (p_g - c_g)g = c_g g + \frac{1 - \alpha_m}{\alpha_m} p_g g$$

It is strictly increasing in p_g at fixed $q > 0$. The allowed minimum $p_g = c_g$ therefore attains the fiber minimum and gives (109). The owner's income is the difference between gross procurement and the input bill, equal to the stated positive rent. Necessity sales at resource cost give the gate no additional margin; the owner's consumption or saving of its income is already financed by this same procurement account. The construction thus counts both real input use and the household income that implements production.

Finally, $1/\alpha_m = 1 + 1/\eta$ and $(1 - \alpha_m)/\alpha_m = 1/\eta$. Substituting (110) makes the coefficient $(c_g/\alpha_m)[c_g/(\kappa \alpha_m^2)]^{-1} = \kappa \alpha_m = \kappa/(1 + 1/\eta)$, proving the normalized form. The zero-output outcome is the separate exit option, or the limit at a zero procurement quote when that quote is included in the feasible set. $\square$

The no-subsidy condition is part of this result. Allowing p_g to approach zero would change the minimum and can remove its attainment. Likewise, endogenous household labor, flexible necessity margins or changing local prices can change the homogeneous envelope. The general regional problem therefore retains $\mathcal{C}(q)$ as an attained correspondence-based object; the power branch is an explicit economic specialization used when closed world price formulas are desired.

A.5 The no-mining corner of the distress benchmark

The interior expression in Proposition 2.7 is used only when mining labor is positive and the household survives. If wealth allows zero mining, the no-mining allocation solves

$$\max_C \{\log(C - c) + \eta \log(1 - dC) : C > c, \quad dC < 1, \quad BC \leq a\}$$

The unconstrained consumption candidate is $C^0 = (1 + \eta dc)/[d(1 + \eta)]$, found from $(C - c)^{-1} = \eta d/(1 - dC)$. If the budget binds first, use $C = a/B$; the chosen candidate must lie in the strict survival domain. The no-mining allocation is an optimum of the complete benchmark only when the marginal gain from beginning mining is weakly nonpositive at that candidate. This procedure solves the actual corner problem. It avoids extending an interior formula to an inadmissible subsistence boundary or assuming that all available wealth must be spent when leisure is also valuable.

B Resilience: primitives, boundaries and analytical examples

B.1 Spatial arbitrage and genuine reservation orders

The results in this subsection and those that follow prove the conditional institutional comparison in Theorem 3.3. They do not replace the broader, committed-institution resilience results.

Lemma B.1 (The spread and the opportunity increment). *At a retained stock y , the opportunity value of one additional package is $H(y + \bar{q}) - H(y)$. If both endpoint optimizers use laundering alone, it is*

$$H(y + \bar{q}) - H(y) = \bar{q}\theta_E - [V(y + \bar{q}) - V(y)]. \quad (112)$$

On that branch the increment exceeds the acquisition bill $\bar{q}\theta_B$ exactly when $V(y + \bar{q}) - V(y) < \bar{q}(\theta_E - \theta_B)$. For a general joint optimizer the seller orders are $\rho_j(a)$ from (27), including residual stock and channel capacities.

Proof. The feasible alternatives to selling the additional package are the joint plans in (22). Their optimized value difference is therefore the stated increment. Substituting the two pure-laundering endpoint values gives (112); subtracting $\bar{q}\theta_B$ gives the strict inequality. An active storage channel or a binding capacity can change the optimizers, so the pure-laundering formula cannot be imposed on such a branch. For $a = q - n\bar{q}$, retaining j optional packages has value

$$H(a) + \sum_{m=1}^j \rho_m(a) = H(a + j\bar{q}).$$

Thus one joint optimization, with unavoidable stock a included, supplies the whole seller order book. Concavity makes its increments decreasing. No separate maximization of laundering and storage may allocate the same unit twice. The expenditure already paid to acquire q is common to every allocation of that stock; subtracting it again from the reservation orders would change the allocation without changing any available opportunity. $\square$

The spread is gross. For example, let $q = \bar{q} = 1$, $\theta_B = 8$, $\theta_E = 10$, $V(\ell) = 3\ell$ and $S = 0$. Then $H(1) = 7$. With one buyer valued at 7.5, genuine pivot revenue is 7 and the continuation payoff is below the acquisition bill 8. Even a positive spread and a sale do not prove profitable production. Actual pivot payments reflect excluded rivals and joint retention, rather than automatically equaling the world quote.

B.2 A complete trading game and its efficiency restriction

Fix q , $Q = \lfloor q/\bar{q} \rfloor$, buyer capacities and $\tau = 0$. Write $b_i = (b_i^1, \dots, b_i^{\bar{k}_i})$ and

$$B_i = \{b_i : M_i \geq b_i^1 \geq \dots \geq b_i^{\bar{k}_i} \geq 0\}, \quad M_i < \infty.$$

Each B_i may instead be any compact convex subset of this set containing zero. Buyers know their own vectors; these vectors are independently drawn from publicly known laws on their declared domains. Incentive and participation requirements apply to every type, including boundary types. Nonidentical distributions are allowed. Values are private: $v_i(k; b_i, 0) = \sum_{j \leq k} b_i^j$ and $v_i(0) = 0$. Independent draws of the vectors do not require independence of one buyer's marginal values.

The true schedule H and all public primitives are common knowledge before the seller chooses the selling procedure. This is a substantive information restriction. Procurement precedes buyer reports and supplies a fixed stock. The finite set $\mathcal{K}_q$ consists of all integral buyer allocations within their capacities with $\sum_i k_i \leq Q$; the seller retains $y = q - \bar{q} \sum_i k_i$. Its own realized operating value is $H(y)$. Write

$$F_H(k; b) = \sum_i v_i(k_i; b_i, 0) + H(q - \bar{q} \sum_i k_i), \quad S_H(b) = \max_{k \in \mathcal{K}_q} F_H(k; b). \quad (113)$$

This S_H is $S_Q(q, b, 0)$, and $S_{H,-i}$ removes the whole buyer. The physical residual always remains with the seller. Efficiency is relative to this package technology; fractional repackaging is not an omitted trade.

Timing and feasible deviations. The seller publicly commits to an initial procedure, its messages, information disclosure, allocation and monetary payment rules. It may announce a different retention schedule or offer only $n < Q$ packages in that procedure. Buyers then observe their types, choose whether to enter, and play its message game. A buyer can refuse initial participation without paying or receiving mineral. After the initial outcome, every original buyer, including a nonparticipant, remains eligible for

terminal trade. There is no new entrant, new private signal, compulsory type revelation or new outside subsidy. Public auction reports reveal reports; they are not a subsequent independent observation of true values.

Initial ownership is enforceable, but the seller cannot irreversibly destroy, process, export, lock away or encumber unsold packages before this final trading opportunity. The same restriction applies to buyers. After final settlement the owner realizes its specified operating value. The horizon ends there. The seller can retain its stock by declining an unfavorable trade; absence of an irreversible withholding commitment is not forced sale.

A concrete continuation protocol is one terminal multilateral contract. The seller can pass or propose a feasible final allocation and a cash-receipt vector $(\delta_0, \delta_1, \dots, \delta_N)$ with $\sum_{j=0}^N \delta_j = 0$. The proposal specifies all quantities and transfers, which are enforceable only if every buyer gives final consent. Rejection by any buyer leaves the initial ownership and transfers unchanged, and there is no later round. Consent needs only the buyer's own known value and stated transfer. A buyer accepts weak improvements and vetoes a strict loss; the seller can always pass. There is no profitable gift at an already efficient outcome, and passing selects no trade at indifference. These rules permit trades of seller-held stock as well as trades among buyers. They do not give the seller perfect information following a hidden deviation.

More general initial procedures and continuation protocols can be compared under the following explicit definition. All monetary flows have finite expectations, all private transfers balance, and every physical opportunity and real cost is already in the terminal values. We exclude wealth effects, budget constraints, externalities in values, and unmodeled rents from information or participation itself. Utility is quasilinear in money.

Definition B.2 (Admissible efficient complete outcomes). An admissible comparison consists of a publicly committed initial procedure and a Bayesian Nash equilibrium of the entire message, participation and continuation game with the information and feasible set above. At every feasible type, including probability-zero boundary types, the specified strategy is an interim best response. At every type profile its terminal allocation maximizes the true F_H over $\mathcal{K}_q$. Each buyer can secure zero by declining initial participation and all subsequent trades. A unilateral deviation can mimic any other type's full strategy, including its participation and resale actions.

Terminal efficiency is the perfect-resale *outcome restriction* used in the theorem. It is not asserted to follow from this finite contract protocol, absence of destruction, liquidity or voluntary bargaining. Asymmetric information and bargaining power can prevent its fulfillment. In particular, an inefficient equilibrium of an alternative reserve auction is outside this comparison even if other equilibria are efficient. We prove that the comparison class is nonempty by constructing a genuine-pivot equilibrium. We do not assume that equilibrium or its seller payoff in Definition B.2.

Relation to the original sources. [Ausubel and Cramton \[1999\]](#), Sections 4–5 and Theorem 4, use a complete-game reduction and distinguish buyer resale from trade in-

volving seller-held quantity. Their Section 5.2 treats attainment as a separate obligation; Theorem 5 uses a coalition-rationality condition. Their divisible-unit, scalar-signal model with known zero seller value does not itself prove the present nonlinear-retention, vector-type result. Krishna and Perry [1998], Lemma 1, Section 5 and Theorem 1, establish payoff equivalence and optimal generalized VCG payments on their independent, compact convex type domain. An optimal basis need not be a Clarke exclusion basis. The next proofs establish the necessary specialization and extension directly, including the normalization of the paper's actual payments.

B.3 Reduction and complete cash accounting

Lemma B.3 (Direct representation of the complete equilibrium). *Every outcome in Definition B.2 induces an efficient, buyer Bayesian incentive-compatible and interim individually rational direct mechanism. Its expected seller continuation value is*

$$\mathbb{E}[T_0^{\text{net}} + H(y^f)] = \mathbb{E}S_H(b) - \sum_i \mathbb{E}U_i(b_i), \quad (114)$$

where U_i is the buyer's interim utility in the entire game.

Proof. Ask for a vector from each buyer and simulate the equilibrium strategy of the reported type throughout the initial procedure and continuation, using the same public and private randomization. Return the final physical allocation and the cumulative net monetary transfers. A true type can perform this imitation in the original game: all of its own information and available actions are specified, and no outside party automatically reveals its true type. The original equilibrium inequalities therefore imply the direct mechanism's Bayesian incentive inequalities. Independence means a type changing its message integrates against the same law of the other buyers. The opt-out strategy secures zero throughout both stages, proving interim individual rationality. Terminal efficiency is precisely the efficiency restriction in Definition B.2.

For clarity, let p_i^a be initial payment to the seller and δ_i the cumulative resale cash received by buyer i ; the seller receives δ_0 , with $\delta_0 + \sum_i \delta_i = 0$. Define

$$P_i = p_i^a - \delta_i, \quad T_0^{\text{net}} = \sum_i p_i^a + \delta_0 = \sum_i P_i, \quad u_i = v_i(k_i^f; b_i, 0) - P_i.$$

Consequently, state by state,

$$T_0^{\text{net}} + H(y^f) = F_H(k^f; b) - \sum_i u_i = S_H(b) - \sum_i u_i.$$

Taking expectations proves (114). Initial intermediate holdings are not consumed as well as terminal holdings. Seller cash from resale is not discarded from the account. External sales proceeds and channel costs already represented in v_i and H must not also be recorded

as private transfers P_i . In particular, $S(s)$ cannot include receipts from this same terminal contract. These restrictions account for every channel and physical use once. $\square$

B.4 The efficient payment bound with the actual pivot basis

Lemma B.4 (Zero is a most-reluctant type). *For every b_{-i} ,*

$$S_H(0, b_{-i}) = S_{H,-i}(b_{-i}). \quad (115)$$

The current whole-buyer pivot gives interim utility $U_i^V(b_i) = \mathbb{E}_{-i}[S_H(b_i, b_{-i}) - S_{H,-i}(b_{-i})]$. It is nonnegative, minimized at zero, and $U_i^V(0) = 0$.

Proof. Any allocation excluding i is feasible when i is present. Conversely, at type zero, returning all of i 's packages to the seller removes zero buyer value and weakly raises H . Maximization gives equality in both directions in (115). Substitution of the actual pivot payment in (24) gives the utility formula. Every nonnegative type weakly raises welfare above the exclusion value, proving the minimum and the normalization. Whole-buyer removal, including losing marginal orders, is essential to this equation. $\square$

Theorem B.5 (Payoff equivalence and the efficient seller-value bound). *For any efficient buyer-BIC, interim-IR direct mechanism M on the stated domain, there are constants $c_i = U_i^M(0) \geq 0$ such that*

$$U_i^M(b_i) = U_i^V(b_i) + c_i, \quad (116)$$

$$\mathbb{E}R^M = W_Q(q, 0) - \sum_i c_i \leq W_Q(q, 0). \quad (117)$$

Here R^M includes cash and true terminal retention. If M uses the same efficient allocation selection as pivot, its interim expected payment is the pivot payment minus c_i . The seller-value bound does not require identical selections at efficient ties.

Proof. For a quantity k_i define $x_i(k_i) \in \{0, 1\}^{\bar{k}_i}$ by $x_i^j(k_i) = \mathbf{1}_{\{j \leq k_i\}}$. Buyer value is $b_i \cdot x_i$. Fix a vector $t \in B_i$ and the ray $b_i = rt$, $0 \leq r \leq 1$. Convexity and inclusion of zero keep this ray in B_i . Let

$$u(r) = U_i^M(rt), \quad g(r) = \mathbb{E}_{-i}S_H(rt, b_{-i}), \quad a(r) = t \cdot \mathbb{E}_{-i}x_i(k_i^M(rt, b_{-i})).$$

Bayesian IC at types rt and st yields

$$u(s) \geq u(r) + (s - r)a(r), \quad u(s) \leq u(r) + (s - r)a(s).$$

Thus u is convex and Lipschitz, with $a(r)$ a subgradient. The bound on x_i supplies a finite Lipschitz constant independent of transfers. For g , evaluate the welfare problem at the allocation selected by M at rt . Efficiency gives $g(s) \geq g(r) + (s - r)a(r)$, with the reverse comparison after exchanging r, s . Hence g is also convex and Lipschitz with

that subgradient. At almost every r both derivatives exist and equal $a(r)$. Absolute continuity gives $u(1) - u(0) = g(1) - g(0)$. This argument allows randomized efficient selections; scalar slopes coincide wherever the convex potential is differentiable, so ties cause no missing term.

Lemma B.4 implies $g(0) = \mathbb{E}_{-i} S_{H,-i}$ and $g(1) - g(0) = U_i^V(t)$. The type t was arbitrary, proving (116); IR at zero gives $c_i \geq 0$. If the terminal allocation selection is the same, subtracting the two utility identities from the same expected operating value gives the claimed interim payment difference. For different efficient tie selections use the full accounting identity instead:

$$\mathbb{E}R^M = \mathbb{E}S_H - \sum_i \mathbb{E}U_i^V - \sum_i c_i = \mathbb{E}R_Q^V - \sum_i c_i.$$

The second equality is the same pointwise accounting for the original pivot rule. By definition $\mathbb{E}R_Q^V = W_Q(q, 0)$, proving (117). This proves optimality of the exclusion basis here without adding participation fees or replacing the mechanism by another generalized VCG. $\square$

The zero-type requirement has content. With one package, $H = 0$ and a single buyer uniform on $[1, 2]$, the efficient allocation always sells. Pivot charges zero. The efficient mechanism that always sells and charges 1 is IC and IR, and yields strictly greater revenue. Its lowest type has zero rent, whereas pivot leaves that type rent 1. Thus the general efficient-payment theorem cannot justify the current pivot on every possible domain of private demand.

B.5 Attainment against auction and resale deviations

Lemma B.6 (Complete-game attainment). *For the genuine- H , all-offered pivot auction followed by the terminal consent protocol in Section B.2, truthful buyer reports and no terminal trade form an ex post Nash equilibrium of the fixed-procedure game. No buyer profits by combining an initial misreport or nonparticipation with any terminal-trade deviation. The equilibrium attains $W_Q(q, 0)$.*

Proof. Consider the strategy profile in which buyers report truthfully, each buyer consents precisely to weak improvements over its initial endowment, and the seller passes after every initial report history. Hold all other reports truthful. Let a be the allocation after any report by buyer i , and let y_a be its retained stock. The pivot formula gives

$$p_i^a + H(y_a) + \sum_{j \neq i} v_j(a_j; b_j, 0) = S_{H,-i}(b_{-i}). \quad (118)$$

The left side is precisely the aggregate initial payoff of the seller and all other buyers: their mutual auction transfers cancel. It is independent of i 's report. Their stage-one allocation also maximizes their true combined value conditional on a_i ; improving it holding a_i fixed would contradict auction efficiency at the reported profile.

Let s_j be each participant's gain from the terminal contract relative to its initial ownership, including its terminal cash receipt. For the seller, $s_0 = H(y_f) - H(y_a) + \delta_0$. For a buyer, $s_j = v_j(k_j^f; b_j, 0) - v_j(a_j; b_j, 0) + \delta_j$. With final veto against the initial endowment, every nondeviating buyer has $s_j \geq 0$; the seller has $s_0 \geq 0$ since it can pass and evaluates its own true H . Under the specified no-trade seller strategy this also holds at every deviation history. More generally the argument is valid for any continuation protecting the nondeviators' initial payoffs. Closed cash accounts give

$$\sum_{j=0}^N s_j = F_H(k^f; b) - F_H(a; b) \leq S_H(b) - F_H(a; b).$$

It follows that

$$v_i(a_i; b_i, 0) - p_i^a + s_i \leq S_H(b) - S_{H,-i}(b_{-i}), \quad (119)$$

which is i 's truthful pivot utility. This is a pointwise argument, so randomized deviations cannot improve it. Efficiency of the final allocation *after a deviation* was not assumed.

If i initially declines to participate, the auction without i gives the complement value $S_{H,-i}$, with $a_i = 0$ and $p_i^a = 0$. Equations (118)–(119) apply unchanged. This checks the ability to skip the auction and join resale, rather than substituting an artificially closed resale market for participation.

At the truthful profile the initial allocation is efficient. Any final contract to which the buyers consent leaves each of them weakly better off. Total welfare cannot increase, so the seller cannot gain strictly from proposing such a contract. Passing is optimal. Buyer vetoes of strict terminal losses are optimal because rejection ends trade. Together with (119), these strategies establish an ex post Nash equilibrium, and therefore a Bayesian Nash equilibrium under the stated prior. The original receipts and retention are realized with no subsequent seller transfers, yielding $W_Q(q, 0)$.

This establishes existence for the complete fixed-procedure game. It does not prove sequential rationality at every off-equilibrium information set, uniqueness, or convergence of an arbitrary bargaining process. The seller's earlier choice among procedures is addressed by the next comparison, not by buyer dominant-strategy truthfulness. $\square$

B.6 Institution choice and the genuine-announcement corollary

Proof of Theorem 3.3. Lemma B.3 represents any admissible efficient complete outcome as an efficient BIC and IR direct mechanism with the exact seller account. Theorem B.5 bounds its expected continuation payoff by $W_Q(q, 0)$. Lemma B.6 constructs an equilibrium within this class that attains the bound with the actual whole-buyer pivot payments and genuine joint retention. This proves maximum expected total seller payoff over the stated class, not pointwise revenue dominance at every bid vector or optimality over inefficient complete outcomes. $\square$

Proof of Corollary 3.4. For every admissible announcement $\widehat{H}$, evaluate buyers' actual complete-game equilibrium, all subsequent seller cash flows, and terminal inventory at the true H . That complete outcome is one of those bounded in Theorem 3.3. At the genuine announcement select the attaining equilibrium in Lemma B.6. Thus

$$\Pi_q(\widehat{H}; H) \leq W_Q(q, 0) = \Pi_q(H; H).$$

The inequality holds for every admissible equilibrium of each alternative; the notation does not promise a unique equilibrium payoff if several exist. In particular, evaluating a distorted reserve against mechanically truthful initial bids while ignoring subsequent behavior is not this comparison. No strictness is implied. Adding a constant to all reported retention values changes neither allocations nor pivot differences, and different reports can agree on every opportunity increment reached on the supported domain. $H(0) = 0$ may be imposed as a harmless normalization when the primitives have that property, but it is not a uniqueness proof. $\square$

B.7 Private seller information and failure of the broad claim

The public- H result is not a truthfulness theorem for a privately informed cartel. A mechanism announcement can signal the seller's type. Buyers then condition their actions and participation on that signal; conditional incentive constraints cannot be inferred by holding an unobserved H fixed. Even independently drawn seller and buyer types do not remove this obstacle.

Proposition B.7 (A private-retention impossibility within the model). *Let $q = \bar{q} = 1$, one buyer have value $b \sim U[0, 1]$, and the seller privately know $c \sim U[0, 1]$, independently, with $H_c(y) = cy$. No internally balanced mechanism can be efficient, Bayesian IC for both parties, and interim IR relative to no purchase for the buyer and retention for the seller. Hence a general efficient truthful-announcement theorem for privately known H is false even within this simple specialization.*

Proof. Efficiency requires trade $x(b, c) = \mathbf{1}_{\{b \geq c\}}$ up to zero-probability ties. Let $U_B(b)$ be buyer interim utility and $U_S(c)$ the seller's interim gain relative to retaining its package, so seller realized gain is net receipts minus cx . The one-dimensional incentive inequalities, integrated as in Theorem B.5, imply

$$U'_B(b) = \Pr(c \leq b) = b, \quad U'_S(c) = -\Pr(b \geq c) = -(1 - c)$$

almost everywhere. IR gives $U_B(0) \geq 0$ and $U_S(1) \geq 0$, whence

$$U_B(b) \geq b^2/2, \quad U_S(c) \geq (1 - c)^2/2.$$

Integrating the two lower bounds gives expected total trading rents at least $1/3$. Internal balance, however, gives their exact sum as

$$\mathbb{E}[(b - c)x] = \int_0^1 \int_c^1 (b - c) db dc = 1/6,$$

a contradiction. A minimum expected outside subsidy of $1/6$ would be needed just to cover this rent shortfall under efficiency. The schedule H_c is compatible with the joint retention definition: for example set $S = 0$, $\theta_E = 2$, $V_c(\ell) = (2 - c)\ell$ and unrestricted laundering. A procurement quote $\theta_B = 1 < 2$ preserves the spatial spread; it is sunk in this fixed-stock example. Thus the spatial wedge does not cure the information problem. $\square$

If a seller's initially private retention schedule is independently and costlessly verified and publicly disclosed *before* mechanism choice, and the conditional buyer law has the stated product structure, the public theorem applies at each verified schedule. This is optimal choice after verification, not incentive compatibility of unverifiable self-reporting. Verification costs, if any, require a separate procurement account. With hidden H independent of buyers, the buyer-envelope proof can also be run after averaging S_H over H , giving an ex ante upper bound $\mathbb{E}_H W_Q(q, 0; H)$. It supplies neither typewise attainment nor seller IC, as the preceding counterexample demonstrates. Correlated buyer types or conditioning events that destroy their independence require a different payment analysis; the manuscript's arbitrary-law resilience identities remain statements within the fixed mechanism.

B.8 Offered capacity, procurement and selective compliance

Proposition B.8 (Scope extensions that preserve the same continuation). *The following extensions hold under the additional conditions stated here.*

1. *At fixed q , an initial offered-capacity restriction $n < Q$ cannot beat $W_Q(q, 0)$ if all unsold packages remain tradable and its complete equilibrium is admissible under Definition B.2.*
2. *If the same comparison and attainment hold at every feasible stock, under a common buyer population and law, publicly known retention technology and common information before procurement, optimal procurement solves $\max_q \{W_Q(q, 0) - \mathcal{C}(q)\}$ whenever that maximum is attained.*
3. *At a fixed $\tau \geq 0$ the theorem extends with the unchanged pivot rule if the compliance burden is determined solely by the final holder and final quantity, cannot be bypassed through an intermediary, and is counted exactly once in v_i . The type domain is the same raw nonnegative marginal-vector domain; final net marginal values may be negative.*

Proof. For the first assertion, the initial capacity changes the initial procedure, not $\mathcal{K}_q$ or H in the terminal comparison. Reduction and the payment bound therefore apply

unchanged. If an unoffered package is permanently unavailable, this is a different feasible set and the assertion does not apply. In particular the static W_n of a committed restricted auction need not equal its continuation value once resale is added.

For the second, at each q every admissible procedure yields at most $W_Q(q, 0)$, and genuine pivot attains it. Subtract the same attained regional implementation cost $\mathcal{C}(q)$ and maximize over q , including exit. At any maximizing stock the mechanism and regional cost-minimizing implementation attain the overall optimum. Randomizing before learning buyer types cannot exceed the supremum of these net values. Compactness, upper semi-continuity or the existing regional attainment hypotheses are still needed for existence; fixed-stock optimality does not prove them.

For the third, put $t_i = \bar{q}\tau \mathbf{1}_{\{i \in \mathcal{R}\}}$ and use $v_i(k; b_i, \tau) = b_i \cdot x_i(k) - t_i k$. The public compliance term is an intercept in the welfare potential along any type ray. The envelope proof is unchanged. At raw type zero the buyer's net value is nonpositive; returning its packages to the seller weakly increases welfare, so (115) still holds and the pivot normalization needs no extra fee. Budget-balanced private transfers and true terminal values give the same accounting and combined-deviation proof. If compliance is a payment to a passive authority rather than a resource cost, the bound is on private surplus net of that payment; tax revenue is not a seller receipt. $\square$

A direct-auction acquisition charge generally fails this final-incidence condition. For a transparent one-package example, let a regulated buyer's gross operating value be 10, an unaffected intermediary's own use value be 0, and the direct acquisition charge be 4. With $H = 0$, direct regulated acquisition has net operating surplus 6. If the intermediary may obtain the package without the charge and transfer it tax-free to the regulated buyer for 3, both can gain relative to the direct allocation with zero auction transfers: the regulated buyer gets 7 instead of 6 and the intermediary 3 instead of zero, producing private surplus 10. The static net value $10 - 4$ no longer describes that route. Anticipating this route also gives the intermediary a resale opportunity; its actual auction behavior cannot be calculated using a zero valuation. The example identifies a missing incidence primitive, not an equilibrium calculation with those inconsistent bids. The original selective-policy exposure identities remain valid within their stated committed auction; they become optimal-institution comparisons only under the no-bypass conditions just proved.

B.9 Why liquidity alone does not select the reserve

There are two additional boundaries. First, at a realized one-package profile with true retention 4 and buyer values 10, 6, genuine pivot collects 6; announcing retention 8 instead collects 8 in the static auction. Buyer truthfulness in each committed auction does not make the seller dominant-strategy truthful. This statewise comparison is not a comparison of admissible complete equilibria across the whole type domain.

Second, let one buyer have $b \sim U[0, 1]$, $H = 0$ and one package. If the seller can withhold stock from all later seller-buyer trade, resale only among buyers is vacuous.

Genuine pivot sells at zero. A posted reserve r earns $r(1 - r)$, maximized at $r = 1/2$ with revenue $1/4$. This is a complete counterexample to optimality from buyer-only resale: it leaves gains from seller-stock trade unrealized when $0 < b < 1/2$. Even if stock remains physically available, the one-shot contract protocol need not remove such bargaining distortions. Requiring efficient terminal outcomes is therefore a real restriction on the comparison, not an empirical conclusion from durability, a deep gold market, or a positive regional/world price spread. Resale costs, price impact, delayed trade and mineral-specific access restrictions can further change the feasible surplus and require their own model. No unproved approximate-optimality bound is inferred from small frictions.

Finally, the incentive claim concerns unilateral deviations of fixed buyer identities, not bidding cartels. In the one-package profile 10,6 with true retention 4, the high buyer pays 6 and earns 4. If the two buyers coordinate so the low buyer reports zero, the high buyer pays 4 and earns 6. A side payment of 1 from the high buyer to the low buyer gives both a strict gain over truthful separate bidding. Seller receipts fall even though the final allocation is efficient. Thus no group-strategyproofness, coalition-proof equilibrium, resistance to false identities or optimality against endogenous buyer cartels follows from the unilateral BIC comparison. These require different participation, information and deviation sets.

B.10 Primitive operating values and joint retention

Concavity of a buyer's operating value gives decreasing marginal orders by the midpoint inequality $2J_i^0(k\bar{q}) \geq J_i^0((k-1)\bar{q}) + J_i^0((k+1)\bar{q})$. Subtracting consecutive increments gives $b_i^k \geq b_i^{k+1}$. The policy subtracts the same $\bar{q}\tau$ from every regulated increment, so it preserves decreasing orders within that buyer. A jointly concave processing return net of convex financing and operating costs, optimized over a convex opportunity set, supplies one primitive interpretation. Mixing two feasible operating plans then yields a feasible plan at the mixed mineral quantity and at least their weighted net return.

For the outside-source bound, take any plan using an additional $\bar{q}$ auction units. If identical units can be obtained outside the auction at delivered price θ_E with no financing or capacity restriction, the plan can be reproduced by purchasing those units outside. Thus $J_i^0(x) \geq J_i^0(x + \bar{q}) - \bar{q}\theta_E$, proving $b_i^k \leq \bar{q}\theta_E$. The alternative must be included in the buyer's operating opportunities. An arbitrary gross processing revenue function would not establish this bound.

For each fixed stock y , the feasible channel allocations in (22) lie in a compact simplex when channel capacities are closed. Continuous channel values attain the maximum. If (ℓ_1, s_1, d_1) and (ℓ_2, s_2, d_2) are optimal at two stocks, their convex mixture is feasible at the corresponding mixed stock. Convex V and concave S imply

$$H(\lambda y_1 + (1 - \lambda)y_2) \geq \lambda H(y_1) + (1 - \lambda)H(y_2)$$

Adding new mineral to free disposal retains the old maximizing plan, so H is nondecreasing. Continuity and finite value on the economically relevant stock domain are maintained. Full packages and physical residuals remain in this single stock constraint throughout.

B.11 Individual opportunities and total seller value

Let y be the actual retained stock and $d_i = k_i^* \bar{q}$ the physical quantity sold to buyer i . The exact aggregate decomposition is

$$R_n - H(q) = \sum_i [p_i - D_i(y)] + \left\{ \sum_i D_i(y) - [H(q) - H(y)] \right\} \quad (120)$$

The opportunity-floor proposition makes the first sum nonnegative. For the second, decreasing increments of concave H give, for $s, t \geq 0$,

$$H(y + s + t) - H(y + s) \leq H(y + t) - H(y)$$

Apply this inequality successively to all blocks d_i to obtain $H(q) - H(y) \leq \sum_i D_i(y)$. This proves nonnegativity of the braces. If H is strictly concave, summing separately evaluated return opportunities can strictly exceed the value of jointly returning all blocks. That is why the aggregate premium need not equal the sum of individual auction premia alone.

The opportunity price also differs from the procurement price. At the allocation stage, every alternative uses the same already purchased q . Moving a block between a buyer and retention compares only their gross continuation values. Subtracting $\theta_B q$ from each alternative cancels it; it cannot change the efficient allocation. Procurement becomes relevant again when comparing different inventories through $W_n(q, \tau) - \mathcal{C}(q)$.

B.12 The original package-depth condition with buyer exclusion

Suppose a residual mechanism offers $K \geq 1$ packages and has no further retention value. There are $u = |\mathcal{U}| \geq 2$ unaffected buyers with decreasing marginal orders. The same calculation is a conservative lower bound when incremental retention opportunities are added. Write

$$K = (u - 1)a + r, \quad a = \lfloor K/(u - 1) \rfloor, \quad 0 \leq r < u - 1 \quad (121)$$

For an order depth j , let $b_{\mathcal{U},[s]}^j$ be the s th highest number among the j th marginal values of unaffected buyers. Missing orders are assigned zero. Define

$$\beta_{\mathcal{U}} = \begin{cases} \min_{h \in \mathcal{U}} b_h^a, & r = 0, \\ b_{\mathcal{U},[r+1]}^1, & a = 0, \\ \min\{\min_{h \in \mathcal{U}} b_h^a, b_{\mathcal{U},[r+1]}^{a+1}\}, & a \geq 1, r > 0. \end{cases} \quad (122)$$

Proposition B.9 (A sufficient unaffected-demand depth). *For any fixed information law and acquisition quote,*

$$\mathbb{E}_\nu \beta_{\mathcal{U}} > \bar{q} \theta_B \implies \mathbb{E}_\nu e_{\mathcal{U}}(K, \omega) > \bar{q} \theta_B$$

A simpler sufficient premise is $\mathbb{E}_\nu \min_{h \in \mathcal{U}} b_h^{\lceil K/(u-1) \rceil} > \bar{q} \theta_B$. These are expected placement-margin conditions on the fixed population, not global production conditions.

Proof. Remove any possible winner. If $r = 0$, its $u - 1$ unaffected rivals each supply a orders no lower than $\min_h b_h^a$, yielding K competing orders. If $r > 0$, those orders supply $(u - 1)a$. Among the $r + 1$ buyers with the largest $(a + 1)$ st orders, at least r remain after any one buyer is excluded. They supply the missing competing orders. The case $a = 0$ uses these first orders alone. Hence $e_{\mathcal{U}} \geq \beta_{\mathcal{U}}$ state by state. Decreasing values imply $\min_h b_h^{\lceil K/(u-1) \rceil} \leq \beta_{\mathcal{U}}$. Taking expectations proves both conclusions. $\square$

The denominator is $u - 1$ because a winning buyer does not compete against itself when its payment is computed. For example, if two unaffected buyers must price four residual packages, each must be able to replace all four units when the other buyer is removed. Their second orders can describe a two-and-two allocation; their fourth orders are needed for this conservative payment floor. A fixed precommitted residual auction is a legitimate special case. The joint-retention mechanism instead uses its complete opportunity book and does not independently choose a laundering quantity before deriving the reserve schedule.

B.13 A moving allocation threshold under uncertainty

Take one package and one regulated buyer with value uniformly distributed on $[0, M]$. Set $q = n = \bar{q} = 1$ and $H(y) = hy$, with $0 < h < M$. The buyer receives the package if $v - \tau > h$; otherwise the seller launders it. Thus

$$L(v, \tau) = \mathbf{1}_{\{v \leq h + \tau\}}, \quad \mathbb{E}L = \frac{h + \tau}{M} \quad \text{for } 0 < h + \tau < M$$

Every realized type has zero local laundering derivative except at its own threshold. Yet $d\mathbb{E}L/d\tau = 1/M$. The difference quotients near the moving threshold do not have the integrable uniform domination that would justify replacing the derivative of the expectation by the expected statewise derivative. The relevant law and the meaning of a zero response must therefore be stated before applying the chain rule to a stochastic allocation.

B.14 Protected regimes under a fixed type law

For fixed offered capacity, a candidate allocation has welfare affine in τ . The full problem and every buyer-exclusion problem must be checked when certifying a regime. Index them by $a \in \{0, 1, \dots, N\}$ and let k_a^0 be their respective unique baseline allocations on

an inventory interval I . If $F_a(k, q, 0)$ is candidate welfare and $m_{\mathcal{R}}(k)$ its regulated count, define

$$T_{\text{cell}} = \inf_{\substack{q \in I, a, k \text{ feasible in } a \\ m_{\mathcal{R}}(k_a^0) > m_{\mathcal{R}}(k)}} \frac{F_a(k_a^0, q, 0) - F_a(k, q, 0)}{\bar{q}[m_{\mathcal{R}}(k_a^0) - m_{\mathcal{R}}(k)]} \quad (123)$$

with the infimum of the empty set defined as infinity. Each candidate with no smaller regulated count loses at least as much welfare as the baseline, so it cannot newly defeat it. A candidate with a smaller count first ties at the displayed ratio. Hence for $\tau < T_{\text{cell}}$ all full and exclusion allocations remain unchanged and their exposure count is unchanged.

This can provide a common-exposure neighborhood for Corollary 3.13. A unique observed allocation alone is insufficient: a buyer-exclusion allocation may switch earlier. Under a full-support type law, positive-probability types can be arbitrarily near a switch, so the uniform statewise threshold can be zero even though expected marginal inventory value is protected. Affine retention gives a useful analytical example of the latter possibility.

Let $H(y) = hy$ with $h > 0$, offer one package on an interval with $q \geq \bar{q}$, and let independent unit-demand values be U, R . The seller's reserve is $h\bar{q}$, independent of inventory. Its exact receipts are

$$R_1(q, \omega, \tau) = hq + [\min\{U, R - \bar{q}\tau\} - h\bar{q}]_+ \quad (124)$$

If a buyer wins, its payment is the larger of the reserve and the other buyer's value. If neither value exceeds the reserve, all stock is retained. Both cases yield the formula. Independence and the tail-integral identity give

$$W_1(q, \tau) = hq + \psi(\tau), \quad \psi(\tau) = \int_{h\bar{q}}^{\infty} [1 - F_U(t)][1 - F_R(t + \bar{q}\tau)] dt \quad (125)$$

The policy loss is common to every inventory within the specified offered- capacity regime. Thus $W_{q\tau} = 0$ there even when $\psi'(\tau) < 0$ and payment-setting identities switch throughout the type population. A global output theorem still compares inventories outside the interval and other offered capacities.

B.15 An exact procurement switch with profitable placement

The following algebraic example separates payment and production conditions. All quantities and costs here specify a downstream implementation example; they do not replace the regional primitives. Packages have size one, retention is zero, there are four independent unit-demand buyers with values $(10 - \tau, 9, 8, 7)$, and integer implementation cost is $\mathcal{C}(Q) = 2.2Q^2$ for $Q \in \{0, 1, 2, 3, 4\}$. Acquisition price can be two, with the remaining positive bill included in implementation cost. The seller offers every acquired package.

The one- and two-package profits are

$$\begin{aligned}\Pi_1(\tau) &= \begin{cases} 6.8, & 0 \leq \tau \leq 1, \\ 7.8 - \tau, & 1 \leq \tau \leq 2, \\ 5.8, & \tau \geq 2, \end{cases} \\ \Pi_2(\tau) &= \begin{cases} 7.2, & 0 \leq \tau \leq 2, \\ 11.2 - 2\tau, & 2 \leq \tau \leq 3, \\ 5.2, & \tau \geq 3. \end{cases}\end{aligned}\tag{126}$$

These formulas follow because each unit-demand winner pays the next losing order. At three packages, revenue is at most 21 and cost is 19.8, so profit is at most 1.2. At four packages every buyer with positive value can be served without displacing another valuable order; revenue is zero. Exit yields zero. Comparing the first two profits therefore gives the unique preferred quantity two below $\tau = 2.7$ and one above it, with a tie at that point.

At the switch, the two-package critical payment is 7.3, far above the acquisition quote two. The unaffected price floor is seven and survives for every tax. Nevertheless the incremental receipts from a second package equal its incremental implementation cost at the switch: $14.6 - 8 = 2.2(4 - 1) = 6.6$. The first allocation/critical-order regime change occurs at tax two; the global procurement switch occurs at 2.7; the positive-placement floor never fails. The three events are distinct.

B.16 A physical-output switch inside one package band

Set package size one, two unit-demand values $(9, 8 - \tau)$, $H(y) = 10y - 2y^2$ and $\mathcal{C}(q) = 2q^2$ on the physical domain $[0, 2]$. The retention function is increasing and concave there. Below one unit only retention is feasible. On $1 \leq q < 2$, the seller's reserve is $\rho(q) = H(q) - H(q - 1) = 12 - 4q \leq 8 < 9$, so the unaffected buyer wins. Its payment is $\max\{8 - \tau, \rho(q)\}$. Adding retained value gives

$$W(q, \tau) = \max\{H(q), 8 - \tau + H(q - 1)\}\tag{127}$$

At two, the seller orders are $(8, 4)$ and the same formula follows by whole-buyer exclusion, under either all-offered trade or one-package withholding. Offering zero packages never improves it.

Subtract costs and complete squares:

$$\begin{aligned}H(q) - 2q^2 &= 6.25 - 4(q - 1.25)^2, \\ 8 - \tau + H(q - 1) - 2q^2 &= 8.25 - \tau - 4(q - 1.75)^2\end{aligned}$$

Below one, profit is less than six. The maximum on the one-package band is therefore attained at $q = 1.75$ for $\tau < 2$, at $q = 1.25$ for $\tau > 2$, and at exactly those two quantities

when $\tau = 2$. Both have $Q = 1$. Their physical residuals are respectively .75 and .25. Hence an unchanged package count does not imply unchanged extraction. At the higher output, the reserve is five and its own payment exposure ends only at tax three. At the lower output the reserve is seven and exposure ends at tax one. The comparison of their complete losses causes the global switch at two, before the high-output branch's own local cutoff.

B.17 Full support with a protected expected marginal value

Let independent U, R be uniform on $[0, 10]$, use unit packages, $H(y) = 7y$ and $\mathcal{C}(q) = 2q^2$ on all $q \geq 0$. For one offered package the formula (125) gives

$$\psi(\tau) = \begin{cases} (3 - \tau)^2(6 + \tau)/600, & 0 \leq \tau < 3, \\ 0, & \tau \geq 3. \end{cases} \quad (128)$$

Indeed its tail integral is $100^{-1} \int_7^{10-\tau} (10-t)(10-t-\tau)dt$. On the one-package band, profit is $6.125 + \psi(\tau) - 2(q - 1.75)^2$, uniquely maximized at 1.75. Below one unit, $7q - 2q^2$ is below five. At two or more units, offering at least two packages makes every active buyer pay the reserve, since it can be served without displacing the other. Receipts are exactly $7q$; the largest profit in that region is six. Withholding one package adds ψ and yields at most $6 + \psi$, still below the optimum. This compares every stock and every allowed offered capacity.

The unique optimal output is thus $q = 1.75$ at every policy level, while profit falls with ψ . For every small $\varepsilon > 0$, however, the event $7 < R < 7 + \varepsilon$ and $U > 7 + \varepsilon$ has positive probability. Its regulated payment setter switches to retention before ε . There is no common positive statewise allocation interval. Expected marginal inventory value remains seven because the premium is inventory-independent, not because every type follows one unchanged order regime.

B.18 Strictly partial local output transmission

Keep the uniform law but use $H(y) = 10y - 2y^2$ on the one-package band. With $\rho(q) = 12 - 4q$, expected receipts are

$$W_1(q, \tau) = H(q) + \int_{\rho(q)}^{10-\tau} (1 - t/10)(1 - (t + \tau)/10)dt \quad (129)$$

when $0 < \rho(q) < 10 - \tau$. Differentiation gives $W_{q\tau} = -4(10 - \rho)/100$. With implementation cost $2q^2$, set $u = 4q - 2 = 10 - \rho$. The regular optimum solves

$$10 - 8q + .04u(u - \tau) = 0, \quad \Pi_{qq} = -8 + .16(2u - \tau)$$

On $1 \leq q < 2$ and $0 \leq \tau \leq 1$, $2 \leq u < 6$, so the Hessian is strictly negative. The derivative at one is positive and at two negative, giving a unique band optimum. At zero policy the relevant root is $u_0 = 25 - 5\sqrt{19}$ and $q_0 = (27 - 5\sqrt{19})/4$. Its transmission coefficient is

$$\gamma = .04u_0 = 1 - \frac{\sqrt{19}}{5} \in (0, 1), \quad \Lambda = (8 - .32u_0)^{-1} > 0$$

Consequently $q_\tau = -\Lambda\gamma$ and the comparator direct levy gives $q_t^{\text{levy}} = -\Lambda$. This is a strictly partial local production response under full-support uncertainty. The statement is restricted to the regular band comparison just derived.

B.19 Why a receipt bound does not bound the output derivative

An explicit counterexample sharpens the distinction. Let the unaffected unit-demand value be nine and the regulated value uniform on $[4, 6]$. Use $H(y) = 10y - 2y^2$ and cost $\mathcal{C}(q) = 6q^2 - 16q + 12$ on an explicitly specified operating domain $[0, 2]$, with separate exit payoff zero. The costs are positive on this domain. The input data describe an illustrative implementation envelope, not a derived village-price system.

On the one-package band the unaffected buyer always wins, and exact gross receipts are $H(q) + \mathbb{E}[R - \tau - \rho(q)]_+$ with $\rho(q) = 12 - 4q$. For $0 \leq \tau < 1/2$, put $a_\tau = (6 + \tau)/4$. Integrating the uniform density gives operating profit

$$\Pi(q, \tau) = \begin{cases} 26q - 8q^2 - 12, & 0 \leq q \leq a_\tau, \\ -4q^2 + (14 - 2\tau)q + (6 + \tau)^2/4 - 12, & a_\tau \leq q \leq 2. \end{cases} \quad (130)$$

Below the one-package band the first expression is simply retention net of cost. At two, the seller's orders $(8, 4)$ give the same expression as the one-package limit by whole-buyer exclusion. The first branch increases up to a_τ and the second is strictly concave. The unique global maximizer over operating choices and exit is $q^* = 1.75 - \tau/4 > a_\tau$, with positive profit. At tax $1/2$ the branches meet at their unique maximum 1.625. That choice has no further exposure and remains optimal at higher taxes by the protected-baseline corollary rebased at $1/2$.

At the zero-policy optimum, the regulated order sets the payment with probability one half. Yet locally

$$\bar{\chi}(q, \tau) = (6 - \tau - \rho(q))/2, \quad \gamma = \partial_q \bar{\chi} = 2, \quad \Pi_{qq} = -8$$

The policy therefore gives $q_\tau = -1/4$, whereas a direct physical levy at the same baseline gives $q_t^{\text{levy}} = -1/8$. Half transmission to expected receipts coexists with twice the output contraction of this comparator. Throughout the regular branch, actual laundering is $q - 1$: its direct policy derivative at fixed stock is zero, while its total derivative includes the changing optimal stock. Both distinctions come from the actual pivot mechanism rather than an arbitrary revenue function.

C Exact geographic aggregation and household reconstruction

This appendix starts from the regional choice $q \in \arg \max \{W_r(q; \Omega) - \mathcal{C}_r(q; A_r)\}$ and its attained implementation correspondence $\Gamma_r(q; A_r)$. It proves the exact aggregation class, its reverse allocation map and its boundaries. Every integral below is over regions, rather than research owners or mineral transactions. Repeated resale does not add terms to these integrals.

C.1 Measure, scale and the cost infimum

Assumption C.1 (A separable regional population). There are finitely many groups. For each group, the region space is a standard Borel space with a finite, completed, atomless regional population measure μ^{reg} , with the group index suppressed, of mass $0 < m < \infty$. Its regional parameters A_r, κ_r are positive, finite and measurable. A common $\eta > 0$ defines

$$s = 1 + 1/\eta, \quad \alpha_m = 1/s = \eta/(1 + \eta), \quad \beta_r = A_r \kappa_r^{-\eta}, \quad 0 < B = \int \beta_r d\mu^{\text{reg}} < \infty. \quad (131)$$

All $q \geq 0$ are feasible, zero-output exit has cost and payoff zero, and each fiber attains the private implementation cost

$$\mathcal{C}_r(q) = \alpha_m \beta_r^{-1/\eta} q^s. \quad (132)$$

Implementations in different regions impose no unrecorded shared resource, credit, coverage or participation constraint.

Atomlessness permits a geographic interpretation and exact entry fractions below. The power-cost identities themselves hold for arbitrary finite measures, including a finite list with unequal regional weights. All almost-everywhere statements refer to the measure actually specified; replacing that measure with a probability measure requires multiplying its expectation by m .

For finite $Q \geq 0$ let

$$\bar{\mathcal{C}}(Q) = \inf \left\{ \int \mathcal{C}_r(q_r) d\mu^{\text{reg}} : q_r \geq 0 \text{ measurable, } \int q_r d\mu^{\text{reg}} = Q \right\}. \quad (133)$$

The integral is an extended nonnegative integral, so the definition permits candidate allocations with infinite cost.

Theorem C.2 (Exact cost, supply and operator-profit aggregation). *Under Assumption C.1, the cost infimum is attained and equals*

$$\bar{\mathcal{C}}(Q) = \alpha_m B^{-1/\eta} Q^s. \quad (134)$$

Its unique minimizer up to null sets is $q_r^Q = \beta_r Q/B$. For a finite common actual net receipt $r_0 \in \mathbb{R}$ with continuation $W_r(q) = r_0 q$, regional optimization has the unique measurable choices

$$\begin{aligned} q_r^* &= \beta_r [r_0]_+^\eta, & \Pi_r^* &= \frac{\beta_r [r_0]_+^{1+\eta}}{1+\eta}, \\ Q^* &= B[r_0]_+^\eta, & \Pi^* &= \frac{B[r_0]_+^{1+\eta}}{1+\eta}. \end{aligned} \quad (135)$$

The aggregate maximum of $r_0 Q - \bar{\mathcal{C}}(Q)$ gives exactly these quantities and profits. The regional output, implementation-cost and profit functions are integrable.

Proof. Set $d\nu = \beta_r d\mu^{\text{reg}}/B$, a probability measure, and $f_r = q_r/\beta_r$. Every feasible output allocation has $\int f d\nu = Q/B$. If its cost is finite, Jensen's inequality for the strictly convex function z^s gives

$$\int \mathcal{C}_r(q_r) d\mu^{\text{reg}} = \alpha_m B \int f_r^s d\nu \geq \alpha_m B (Q/B)^s. \quad (136)$$

An infinite-cost allocation satisfies the same inequality. The candidate $q_r^Q = \beta_r Q/B$ is measurable, has integral Q , and has cost exactly the right-hand side; the infimum is therefore attained. Equality for finite cost holds only when f is constant ν -almost everywhere. Since $\beta_r > 0$ almost everywhere, this is also μ^{reg} -almost-everywhere uniqueness. For $Q = 0$, nonnegativity directly gives the unique zero choice.

For $r_0 > 0$, strict convexity of cost makes the objective strictly concave. Its derivative is $r_0 - \beta_r^{-1/\eta} q^{1/\eta}$, its right derivative at zero is positive, and it tends to minus infinity as $q \rightarrow \infty$. The unique optimum is therefore $q_r^* = \beta_r r_0^\eta$. At that optimum cost is $\alpha_m \beta_r r_0^{1+\eta}$, so profit is the stated remaining fraction. If $r_0 \leq 0$, any positive output has a negative objective; zero is uniquely optimal. Each nonzero quantity, cost or profit is a finite constant times the integrable function β_r . Integration proves the aggregate identities. The same derivative calculation with B replacing β_r proves the final maximization claim. $\square$

Corollary C.3 (What a nonlinear total continuation may use). *Suppose an institution jointly chooses regional outputs and its attained downstream value is $W(Q; \Omega)$, depending on these outputs only through their total. Then*

$$\sup_{(q_r)} \left\{ W \left(\int q_r d\mu^{\text{reg}}; \Omega \right) - \int \mathcal{C}_r(q_r) d\mu^{\text{reg}} \right\} = \sup_{Q \geq 0} \{ W(Q; \Omega) - \bar{\mathcal{C}}(Q) \}. \quad (137)$$

If W is finite valued and upper semicontinuous and $W(Q; \Omega) - \bar{\mathcal{C}}(Q)$ tends to minus infinity as $Q \rightarrow \infty$, and some finite Q has finite value, the right side attains its maximum. Every maximizing Q has the regional reconstruction q_r^Q and, under Proposition C.6, a household reconstruction. This statement does not aggregate independent local nonlinear continuations into a representative monopoly.

Proof. For any regional allocation, Theorem C.2 bounds its cost below by the cost at its total. Conversely, the explicit cost minimizer achieves equality at every finite total. This

proves both inequalities in (137). Coercivity restricts an upper maximizing sequence to a bounded interval; upper semicontinuity gives an attained limit. Implementation is by the already attained fibers and the selection result. $\square$

For a counterexample to an independent-local-monopoly interpretation, take $\eta = 1$, total region mass one and a local continuation $W_r(q) = vq - dq^2/2$, where $v, d > 0$. The unique local supply is $q_r = v\beta_r/(1 + d\beta_r)$. Population I has $\beta_r = 1$ everywhere. Population II has $\beta_r = 1/2$ on half its mass and $\beta_r = 3/2$ on the other half. Both have $B = 1$, but strict concavity of $b \mapsto b/(1 + db)$ gives strictly less total supply in population II. Consequently its total profit differs as well. The distribution of local market power and capacity must remain in the nonlinear world model.

Finite B is also an economic restriction. On $(0, 1]$ with Lebesgue measure and $\beta_r = 1/r$, $B = \infty$. For each $n > 1$, restricting production to $[1/n, 1]$ yields attained cost $\alpha_m(\log n)^{-1/\eta}Q^s$. Hence the unrestricted cost infimum for $Q > 0$ is zero. It cannot be attained because every positive-output allocation has strictly positive cost somewhere on a set of positive measure. Common positive-price supply is infinite. A “zero aggregate cost” formula would conceal this nonattainment. A group with $B = 0$ may instead be defined by zero-capacity regions with infinite cost for positive output; it has feasible total only $Q = 0$.

C.2 Contracts, accounting moments and invariance

Proposition C.4 (Distinct supply and profit moments). *Consider finitely many groups satisfying Assumption C.1. Groups in origin j receive $r_g = \omega_g p_j$ with $p_j \geq 0$ and $\omega_g \geq 0$. The exact physical output and private operator profit are*

$$Q_j = p_j^\eta S_j, \quad \Pi_j = \frac{p_j^{1+\eta}}{1+\eta} H_j, \quad S_j = \sum_{g \in j} B_g \omega_g^\eta, \quad H_j = \sum_{g \in j} B_g \omega_g^{1+\eta}. \quad (138)$$

The total private implementation cost at this decentralized allocation is $\alpha_m p_j^{1+\eta} H_j$. Except when its allocation happens to have a common q_r/β_r , it exceeds the unrestricted minimum cost of the same total computed with $\sum_{g \in j} B_g$.

Proof. Apply (135) group by group and sum. The cost expression follows from the regional cost at the optimum. Strictness follows from the equality condition in (136); with $p_j > 0$ it requires identical wedges in positive-capacity groups. If contract or origin constraints prevent such reallocation, the unrestricted cost is a mathematical lower bound rather than the implementable cost for the specified contracts. $\square$

For example, two equal-capacity groups with $\eta = 1$ and wedges 1 and $1/2$ have $S = 3B_0/2$ and $H = 5B_0/4$. Substituting S for H overstates operator profit by twenty percent. With nonnegative fees $p_j - r_g$, the identity

$$p_j Q_j = \Pi_j + \int \mathcal{C}_r(q_r) d\mu^{\text{reg}} + \sum_{g \in j} (p_j - r_g) Q_g \quad (139)$$

records the fee account separately. That account's expenditures still require their own resource, transfer and reserve entries. Taxes, refunds, fraud-organizer income and collateral losses belong to the actor that actually receives or pays them; only genuinely linear net amounts can be absorbed into r_g . Profit from a separate fraud business is added to its owner's profit once, without adding fictitious extraction to Q_j .

More generally, a measurable regional receipt $r(r)$ gives $Q = \int \beta_r[r(r)]_+^\eta d\mu^{\text{reg}}$ and $\Pi = \int \beta_r[r(r)]_+^{1+\eta} d\mu^{\text{reg}} / (1 + \eta)$ whenever both integrals are finite. For example, additive fees yield $r(r) = p - f(r)$ and a price-dependent active set. Unbounded receipts require the displayed integrability conditions; finite B alone is insufficient.

The formulas pass four exact scale checks. A measurable partition with a common receipt replaces B by the sum of the subgroups' B values. Replication of the population by $\lambda > 0$ replaces $\mu^{\text{reg}}, B, Q, \Pi$ by $\lambda\mu^{\text{reg}}, \lambda B, \lambda Q, \lambda\Pi$, and $\bar{\mathcal{C}}_{\lambda B}(\lambda Q) = \lambda \bar{\mathcal{C}}_B(Q)$. If mineral is measured in units with $q' = uq$ and currency is measured with money' $= v$ money, then

$$p' = \frac{v}{u}p, \quad \beta' = u^{1+\eta}v^{-\eta}\beta, \quad \mathcal{C}'(q') = v\mathcal{C}(q), \quad \Pi' = v\Pi. \quad (140)$$

These identities follow by substitution, with no change in physical or economic allocations. Finally, alternative parameterizations of A_r and κ_r with the same β_r describe the same power envelope; they need not describe the same research ownership or innovation process.

C.3 Measurable and integrable return to households

Pointwise attainment is not enough to integrate an uncountable collection of household choices. The missing measurable and integrable structure is stated here rather than inferred from individual optimality.

Assumption C.5 (A selectable implementation correspondence). At a specified measurable regional output q_r , each selected finite coverage branch can represent the complete regional allocation, terms and transfers in a common Euclidean space $\mathbb{R}^d$ (padding unused coordinates when needed). The nonempty correspondence $F(r) = \Gamma_r(q_r; A_r)$ is closed valued and weakly measurable: $\{r : F(r) \cap O \neq \emptyset\}$ is measurable for each open set $O \subseteq \mathbb{R}^d$. For the coordinates to be aggregated, there is an integrable nonnegative bound $h(r)$ with $\|z\| \leq h(r)$ for all $z \in F(r)$. Regional resource and budget identities hold at every member of $F(r)$, with integrable entries.

Closedness, measurability and the bound are sufficient conditions. They can be established by explicit formulas, as below, or separately for the chosen regional equilibrium

correspondence. Compact positive-output fibers in the finite Selva economy do not by themselves supply uniform bounds across a heterogeneous population. An independently selected allocation may be used instead when it is already known to be measurable and integrable; no bound on all other implementations is then needed.

Proposition C.6 (Attained measurable household reconstruction). *Under Assumption C.5, there is a measurable $z_r \in \Gamma_r(q_r; A_r)$. Its specified coordinates are integrable. Thus the aggregate output allocation of Theorem C.2 admits household choices attaining its cost, and integrating regional budget and resource identities preserves those identities exactly.*

Proof. We give a selection construction. Enumerate a countable dense set $(y_k)_{k \geq 1}$ in $\mathbb{R}^d$. For each fixed k , the distance $d_k(r) = \inf_{z \in F(r)} \|y_k - z\|$ is measurable: its strict sublevel at $a > 0$ is $\{r : F(r) \cap B(y_k, a) \neq \emptyset\}$. Choose $z_1(r)$ as the first y_k with $d_k(r) < 1/2$. Recursively choose $z_n(r)$, $n \geq 2$, as the first y_k satisfying

$$d_k(r) < 2^{-n}, \quad \|y_k - z_{n-1}(r)\| < 3 \cdot 2^{-n}. \quad (141)$$

Each such first-index choice is measurable, because its countable candidate sets are measurable. The recursion is feasible: since $\text{dist}(z_{n-1}, F) < 2^{-(n-1)}$, a point of F within that distance can be approximated by a dense-set point within 2^{-n} , using strictly smaller errors if necessary. The bounds in (141) make $z_n(r)$ Cauchy. Its pointwise limit $z(r)$ is measurable and has zero distance from the closed set $F(r)$, so it belongs to $F(r)$. The integrable bound then gives integrability of the required coordinates. Integrating finite sums and differences of integrable regional accounts preserves their equalities. Attainment of the cost follows from membership in Γ_r . $\square$

Two failures show why the qualifications matter. If a required resource coordinate is $g_r = 1/r$ on $(0, 1]$, every regional allocation may be finite and attained while its aggregate resource requirement is infinite. In the power construction, choosing $\beta_r = 1$ and $c_r^g = r$ produces exactly this problem at a positive common receipt, up to a fixed factor. Finite aggregate expenditure $\int c_r^g g_r d\mu^{\text{reg}}$ does not imply finite physical input $\int g_r d\mu^{\text{reg}}$. Likewise, with positive consumption surplus $C_r - c_r = \exp(-1/r)$, every household survives and aggregate consumption is finite, but aggregate log utility is minus infinity. An integrable allocation is not automatically integrable welfare.

Proposition C.7 (A fully constructive household branch). *Let Assumption C.1 hold and set*

$$a_r = \left(\frac{c_r^g}{\kappa_r \alpha_m^2} \right)^{\alpha_m} A_r^{1-\alpha_m}, \quad q = a_r g^{\alpha_m}. \quad (142)$$

In each region a finite, fixed number of households own shares $\xi_{hr} \geq 0$ summing to one in this firm. The firm's variable input has resource price $c_r^g > 0$. The gate may post $p_r^g \geq c_r^g$, posts a mineral procurement quote, and sells necessities at resource cost; there are no further costs on this branch. Household h has nonmining wealth w_{hr} backed by owned resources, a basket price $\bar{p}_{hr} > 0$, subsistence requirement $c_{hr} > 0$, cash allowance

$b_{hr}^0 > 0$, and cash weight $\nu_{hr} > 0$. No variable household labor is used, so leisure is one. Its objective is

$$\log(C_{hr} - c_{hr}) + \nu_{hr} \log(b_{hr} + b_{hr}^0), \quad b_{hr} \geq 0. \quad (143)$$

Suppose all primitives and ownership shares are measurable. Assume $w_{hr} - \bar{p}_{hr}c_{hr} \geq \epsilon > 0$, the quantities $\bar{p}_{hr}, b_{hr}^0, \nu_{hr}$ are bounded above and away from zero, $\int \sum_h (w_{hr} + c_{hr}) d\mu^{\text{reg}} < \infty$, and $\int (\beta_r/c_r^g + c_r^g) d\mu^{\text{reg}} < \infty$. Write $p_{\max}$ for a common upper bound on the basket prices $\bar{p}_{hr}$.

For every finite group total Q , write $t = Q/B$ and $q_r = \beta_r t$. The following formulas attain the cost envelope and construct measurable, integrable physical allocations, terms, transfers and finite integrable household utilities:

$$\begin{aligned} \mathcal{C}_r &= \alpha_m \beta_r t^s, & p_r^g &= c_r^g, & \theta_r &= \alpha_m t^{1/\eta}, \\ g_r &= \frac{\alpha_m^2 \beta_r}{c_r^g} t^s, & H_r &= (1 - \alpha_m) \mathcal{C}_r, & M_{hr} &= w_{hr} + \xi_{hr} H_r - \bar{p}_{hr} c_{hr}, \\ b_{hr} &= \left[\frac{\nu_{hr} M_{hr} - b_{hr}^0}{1 + \nu_{hr}} \right]_+, & C_{hr} &= c_{hr} + \frac{M_{hr} - b_{hr}}{\bar{p}_{hr}}. \end{aligned} \quad (144)$$

At $Q = 0$ take the zero procurement quote or the separately feasible exit with the same surviving outside allocation. Procurement accounting obeys

$$\mathcal{C}_r = \underbrace{c_r^g g_r}_{\text{real productive input}} + \underbrace{H_r}_{\text{household operating income}}, \quad \bar{p}_{hr} C_{hr} + b_{hr} = w_{hr} + \xi_{hr} H_r. \quad (145)$$

Proof. At a positive quote θ , strict concavity makes the owned firm's unique input maximize $\theta a_r g^{\alpha_m} - p_r^g g$. Its first-order condition is $\theta q = p_r^g g / \alpha_m$. Every shareholder with a positive ownership share values higher operating income strictly, so the profit-maximizing firm choice is consistent with the household choices. To deliver a prescribed $q > 0$, technology fixes $g = (q/a_r)^{1/\alpha_m}$. The gate's net procurement bill, subtracting its input margin, is

$$\theta q - (p_r^g - c_r^g) g = c_r^g g + \frac{1 - \alpha_m}{\alpha_m} p_r^g g. \quad (146)$$

For fixed target output, this is increasing in p_r^g . The feasible minimum $p_r^g = c_r^g$ attains the fiber minimum, giving cost $(c_r^g / \alpha_m) a_r^{-1/\alpha_m} q^s$. Substituting (142) gives exactly $\alpha_m \kappa_r A_r^{-1/\eta} q^s$. At $q_r = \beta_r t$, the expressions for $\mathcal{C}_r$, g_r , θ_r and H_r follow. Their resource and income split follows directly from the firm's input bill.

The remaining household budget is $\bar{p}_{hr}(C_{hr} - c_{hr}) + b_{hr} = M_{hr} > 0$. Both goods are valuable, so the budget binds. Strict concavity gives a unique choice. At an interior cash choice its first-order condition is $\nu_{hr}(M_{hr} - b_{hr}) = b_{hr} + b_{hr}^0$; at a nonpositive proposed cash amount the constrained optimum is $b_{hr} = 0$. These are precisely (144). With $b = 0$ the positive surplus is $M/\bar{p}$; otherwise it is $(M + b^0)/[\bar{p}(1 + \nu)]$. The assumed outside wealth therefore gives strict survival even when output is zero.

All formulas are measurable. Output, cost and operating income are constant multiples of β_r ; physical input is a constant multiple of β_r/c_r^g . The input quote is integrable by its

separate assumption and the procurement quote is constant within the group. Moreover $0 \leq b_{hr} \leq M_{hr}$ and $0 < C_{hr} - c_{hr} \leq M_{hr}/\bar{p}_{hr}$, so the assumed wealth, subsistence and price bounds give integrability. Consumption surplus is bounded below by $\epsilon/[p_{\max}(1+\bar{v})]$ and $b + b^0$ is bounded away from zero. Their logarithms thus have bounded negative parts; $\log^+ y \leq y$ and the preceding integrability give integrable positive parts. Leisure one adds zero to the Selva leisure term. Finally each household budget and the gate's procurement account hold algebraically; the stated ownership shares exhaust operating income exactly. $\square$

The basket can be the original Selva Leontief basket: fix its physical import requirement $\lambda \in (0, 1)$ and pass-through local and imported prices, so $\bar{p} = (1 - \lambda)p^x + \lambda c^z$. Assign quantities $(1 - \lambda)C$ and λC . A competitive background sector or owned outside asset finances w and supplies those goods at the specified resource prices. These are resources in the branch, not new cartel gifts. Changing these opportunities, introducing endogenous mining labor or allowing input subsidies changes the fiber problem. In particular, the no-subsidy restriction is essential to (146).

This branch separates three objects often conflated in aggregation. Private gate cost is $\mathcal{C}_r$. Real productive-input expenditure is $\alpha_m \mathcal{C}_r$; household operating income is the remaining fraction. Household welfare is the utility of its actual consumption and cash, with outside wealth and all local resources still accounted for. Two populations with identical β_r and different outside wealth or ownership shares have the same power supply but different household welfare. The fixed shares used here also do not grant ownership of future patents.

An explicitly specified physical damage account is equally necessary. For example, if regional current load is $D_r = \delta_r g_r + \zeta_r q_r$, with nonnegative measurable coefficients and finite relevant integrals, the constructive allocation gives

$$D = \alpha_m^2 t^s \int \frac{\delta_r \beta_r}{c_r^g} d\mu^{\text{reg}} + t \int \zeta_r \beta_r d\mu^{\text{reg}}. \quad (147)$$

The proof is substitution and integration of nonnegative terms. A fixed site-opening load or a stock of unrepaired forest damage requires its own additional account. Even the two-variable flow example requires moments beyond B ; a claim of conservation must specify which physical criterion is meant and compare that criterion along the reconstructed allocation.

C.4 Entry, finite coverage and subsistence restrictions

Proposition C.8 (Entry-adjusted supply and its exact homogeneous envelope). *Add a measurable fixed private operating cost $f_r \geq 0$, paid only when $q > 0$, to the power cost. Exit remains feasible. At common receipt r_0 , let*

$$V_r = \frac{\beta_r [r_0]_+^{1+\eta}}{1+\eta} - f_r. \quad (148)$$

If $V_r > 0$ the unique output is $\beta_r[r_0]_+^\eta$; if $V_r < 0$ the unique output is zero. At $V_r = 0$ either exit or the positive unconstrained optimum is optimal, when that optimum is positive. A measurable tie rule gives

$$Q = [r_0]_+^\eta \int \beta_r e_r d\mu^{\text{reg}}, \quad \Pi = \int [V_r]_+ d\mu^{\text{reg}}, \quad (149)$$

where $e_r = 1$ on $\{V_r > 0\}$, $e_r = 0$ on $\{V_r < 0\}$ and $e_r \in \{0, 1\}$ on ties. The individual costs, output and net profits are integrable whenever $\int \beta_r d\mu^{\text{reg}} < \infty$ and $\int f_r d\mu^{\text{reg}} < \infty$.

If $\beta_r = \beta_0 > 0$ and $f_r = f > 0$ throughout an atomless population of mass m , define

$$r_E = \left(\frac{(1 + \eta)f}{\beta_0} \right)^{1/(1+\eta)}, \quad q_E = \beta_0 r_E^\eta. \quad (150)$$

The attained aggregate fixed-output cost is

$$\bar{c}^f(Q) = \begin{cases} r_E Q, & 0 \leq Q \leq m q_E, \\ m f + \alpha_m (m \beta_0)^{-1/\eta} Q^s, & Q \geq m q_E. \end{cases} \quad (151)$$

For $0 < Q < m q_E$, exactly mass Q/q_E operates at output q_E per active region. At and above $m q_E$, all regions operate with output Q/m . At receipt r_E , any aggregate output in $[0, m q_E]$ is optimal; at $r_0 > r_E$, output is $m \beta_0 r_0^\eta$, and below r_E it is zero.

Proof. Conditional on operating, subtraction of f_r does not change the unique positive optimum when $r_0 > 0$. Its value is V_r ; compare that value with exit. At nonpositive receipt, positive output has negative profit and zero is optimal. The activity comparisons and a fixed measurable tie rule are measurable. The integrability claims follow from the bounds by constant multiples of β_r and f_r .

For the homogeneous fixed-total problem, let n be the mass of the active set. Jensen's inequality conditional on this set bounds total cost below by

$$F(n) = f n + \alpha_m \beta_0^{-1/\eta} Q^s n^{-1/\eta}, \quad 0 < n \leq m. \quad (152)$$

The bound is achieved when each active region produces Q/n . Atomlessness of a finite measure permits a measurable subset of every mass $n \in [0, m]$, so each such candidate is feasible. For $Q > 0$, $F(n) \rightarrow \infty$ as $n \downarrow 0$, and differentiation gives

$$F'(n) = f - \frac{\beta_0^{-1/\eta}}{1 + \eta} (Q/n)^s.$$

It vanishes exactly at $n = Q/q_E$ and is strictly increasing in n . Consequently $n^* = \min\{m, Q/q_E\}$. At q_E the zero-profit identity $r_E q_E = f + \alpha_m \beta_0^{-1/\eta} q_E^s$ makes the first segment $r_E Q$; all-active substitution gives the second. At $Q = 0$ choose exit everywhere. Both segments agree and have slope r_E at $m q_E$. Maximizing receipt times output less

this envelope gives the final supply correspondence, equivalently obtained by choosing an arbitrary measurable active subset at the individual tie. $\square$

Proposition C.9 (Every heterogeneous entry total has an attained reconstruction). *Retain the atomless population and power variable costs of Assumption C.1, and allow any measurable finite fixed costs $f_r \geq 0$. Define the entry thresholds and integrated operating profit by*

$$\tau_r = \left(\frac{(1+\eta)f_r}{\beta_r} \right)^{1/(1+\eta)}, \quad P(p) = \int \left[\frac{\beta_r p^{1+\eta}}{1+\eta} - f_r \right]_+ d\mu^{\text{reg}}, \quad p \geq 0. \quad (153)$$

For every finite $Q > 0$ there is a unique $p > 0$ with

$$p^\eta \int_{\{\tau_r < p\}} \beta_r d\mu^{\text{reg}} \leq Q \leq p^\eta \int_{\{\tau_r \leq p\}} \beta_r d\mu^{\text{reg}}. \quad (154)$$

Activate all regions with $\tau_r < p$, an appropriate measurable subset of the tie set $\{\tau_r = p\}$, and no others. Each active region produces $\beta_r p^\eta$. This allocation is measurable and attains the exact heterogeneous fixed-total implementation cost

$$\bar{\mathcal{C}}^f(Q) = pQ - P(p) = \sup_{u \geq 0} \{uQ - P(u)\}. \quad (155)$$

Its output, variable cost and paid fixed costs are integrable, even if $\int f_r d\mu^{\text{reg}} = \infty$. At $Q = 0$, exit everywhere attains zero cost. Household reconstruction again uses the feasible active and exit branches; the theorem does not finance an infeasible exit opportunity.

Proof. Use the finite atomless weighted measure $d\nu = \beta_r d\mu^{\text{reg}}$. Let $Q^-(p) = p^\eta \nu\{\tau < p\}$ and $Q^+(p) = p^\eta \nu\{\tau \leq p\}$. These functions are nondecreasing; the first is left continuous on $(0, \infty)$ and the second is right continuous. Both vanish as $p \downarrow 0$, since they are at most Bp^η . As $p \rightarrow \infty$, the sets $\{\tau < p\}$ increase to the full population up to null sets, so $Q^-(p) \rightarrow \infty$. For a fixed $Q > 0$, let $p = \inf\{u > 0 : Q^+(u) \geq Q\}$. The preceding bounds give $0 < p < \infty$. Right continuity gives $Q \leq Q^+(p)$, while approaching p from below gives $Q^-(p) \leq Q$. Atomless divisibility of ν provides a measurable tie subset with weighted mass $Q/p^\eta - \nu\{\tau < p\}$. This constructs the specified output.

If $0 < p_1 < p_2$ both supported the same positive total, every region selected at p_1 , including its selected ties, would be strictly active at p_2 . Hence $Q^-(p_2) \geq (p_2/p_1)^\eta Q > Q$, a contradiction. The supporting price is unique. At every selected active region, $f_r \leq \beta_r p^{1+\eta}/(1+\eta)$; its variable cost is $\alpha_m \beta_r p^{1+\eta}$. Their integrals are finite because B is finite. The nonnegative profit integrand is bounded by $\beta_r p^{1+\eta}/(1+\eta)$ as well.

For any nonnegative output y and any $u \geq 0$, individual profit maximization with exit implies

$$uy - \mathcal{C}_r(y) - f_r \mathbf{1}_{\{y > 0\}} \leq \left[\frac{\beta_r u^{1+\eta}}{1+\eta} - f_r \right]_+.$$

Integrate this inequality at any feasible allocation of total Q to obtain the lower bound $uQ - P(u)$ on its cost, permitting infinite cost. The constructed allocation at p is individ-

ually optimal in every region, so equality holds pointwise and after integration. It attains both the aggregate infimum and the stated scalar dual supremum. For $Q = 0$, only zero output almost everywhere is needed and its cost is zero. $\square$

Equation (151) is a corrected exact result, not a power cost through entry. For example, with $\eta = \beta_0 = m = 1$ and $f = 1/2$, it is Q for $Q \leq 1$ and $(1 + Q^2)/2$ for $Q \geq 1$. The no-entry power formula $Q^2/2$ is wrong on both ranges. Fractional aggregate entry uses different complete regions; it does not ask one government or one household to pay a fractional setup cost. Fixed costs that are transfers, real installation expenditure or collateral contributions must be classified in the corresponding separate accounts.

For completeness, a finite measurable coverage menu can be handled without extrapolating a power envelope. In branch $a \in \{1, \dots, K\}$ let $0 \leq L_r^a \leq U_r^a < \infty$, $\beta_r^a > 0$, and $f_r^a \geq 0$ be measurable. Conditional on that branch the allowed outputs are the closed interval $[L_r^a, U_r^a]$ and cost is $f_r^a + \alpha_m(\beta_r^a)^{-1/\eta} q^s$. Include the feasible exit separately. Strict concavity gives the explicit candidate

$$q_r^a = \min\{U_r^a, \max\{L_r^a, \beta_r^a[r_0]_+^\eta\}\}. \quad (156)$$

Evaluate its actual net value, choose the best branch or exit, and break ties by a fixed ordering. This produces a measurable attained regional choice, since finite maxima and the displayed formulas are measurable. If the chosen output, costs and household implementations admit integrable bounds, summing them gives an attained, feasible aggregate supply. Different branches generally have different Γ_r^a and different physical damage. Internal village complementarity must be resolved in these regional branches before regions are integrated.

This finite-branch formula covers an output floor required to finance a subsistence deficit, a capacity ceiling, and finite coverage comparisons when their regional envelopes have the stated form. For example, if $q \in \{0\} \cup [L, \infty)$ with $L > 0$ and $\eta = \beta = 1$, then for $0 < r_0 < L$ the operating candidate is L , not r_0 ; with zero fixed cost it operates only when $r_0 L - L^2/2 \geq 0$. If exclusion from coverage removes the outside survival option, zero-output exit cannot be inserted without financing household support. Likewise, a strict subsistence boundary such as $C > c$ is not a compact feasible constraint: its feasible set may be empty or a limiting cost may fail to attain. The general Selva attainment theorem must then be applied with a feasible household allocation. The stated supply formula applies only to regions whose production and subsistence constraints can be satisfied.

D Common-price equilibrium, composition and four accounts

D.1 A complete timing and ownership specification

Fix a date and a public state containing every existing physical holding, origin, cohort, contract, evidence record and unresolved liability. Institutions post the invoice-indexed contracts, their funded capacities, the origin assignment rule and the mechanisms for current stock before private reports. Regional suppliers choose production using the actual quotes. Organizers are competitive in honest procurement and obtain separate compensation for service and bond capital from the fee schedule; each has an exclusive certification slot and the residual profit from fraud in that slot. Entry into these licenses is fixed in this class. The authority adjusts the mass of licensed slots to actual procurement. There is no auction for the fraud franchise whose missing license payment would alter the reported organizer profit.

The slot population is atomless within each origin. The model uses an exact law of large numbers as an explicit population assumption: for every measurable slot subset used by a contract, the realized cost and independent signal fractions equal their stipulated probabilities almost surely. Alternatively all equations are ex ante expected quantities with state-contingent balancing contracts. Ordinary product independence on an unspecified continuum is not by itself a proof of pathwise aggregation. The exact population assumption, rather than unproved elimination of a common shock, is used in the deterministic path construction.

Buyers commit invoices to segregated escrow; funded providers accept contracts; regional commodity deliveries settle against escrow. A fraudulent organizer purchases one unit of true illegal-origin stock at p_D , replacing prospective honest procurement at h . It has no claim on unproduced clean output and cannot source that unit from an origin with zero illegal stock. Procurement of dirty stock is sunk at the subsequent evidence and reporting stage. Comparing honesty and fraud before the purchase gives the spread $h - p_D$; comparing a later diversion of existing stock cancels its purchase payment and uses the originally committed residual sale value. These comparisons are consistent and must not be added together.

The common route commits every remaining unit to sale at p_D under a verified stock contingency. Thus $\mathcal{W}_\Sigma(q - d \mid q, \Omega) = p_D(q - d)$ on this branch. In the nonlinear chapter the corresponding residual mechanism is retained instead. Fees pay the actual ordinary service wages, honest inspection, deployment, installation, maintenance, evidence retention, claims administration and capital expenses, including advance-facility interest and fees; the dynamic funding theorem supplies sufficient bounds. Fines are not relied on to balance honest operation. Buyers face the specified information timing: they cannot sell suppression commitments before obtaining evidence, can bargain with all original partic-

ipants after obtaining it, and optimize all feasible informed mitigation before the printed silence loss is used.

D.2 Fraud and diligence from individual choices

Normalize money by the actual commodity invoice h . Write $A_k = r_k l_{Lk} - l_{Ek} > 0$. A buyer with posterior π minimizes $\kappa_k e^2/2 + \pi r_k l_{Lk} - \pi A_k e$ on $[0, 1]$. Its unique best response is $e = \min\{a_k \pi, 1\}$. At an evidence-stage silence bargain its reservation payment is hA_k , whereas all original participants jointly can avoid at most $h(1 - r_k)b_k$. The strict no-silence inequality rules out a mutually acceptable bargain under the specified accessible side contracts. Reports occur also following off-path genuine evidence. False messages cannot trigger an evidence claim.

Given effort e , the original coalition's required expected compensation plus the organizer's own expected forfeiture is $hb_k[r_k + (1 - r_k)e]$. Each signer receives its expected loss and keeps its ordinary honest wage; an advance facility provides its full possible bribe before the draw. Separate deposits fund collectible liabilities. The organizer chooses fraud exactly when

$$z < 1 - x - b_k[r_k + (1 - r_k)e], \quad z \sim U[0, c_k]. \quad (157)$$

The right side includes the dirty receipts forgone through false admission. Bayes consistency gives

$$\pi = \left[\frac{1 - x - b_k[r_k + (1 - r_k) \min\{a_k \pi, 1\}]}{c_k} \right]_0^1. \quad (158)$$

The right side is continuous and nonincreasing in π ; subtracting π makes it strictly decreasing from a nonnegative to a nonpositive value. There is exactly one solution. The interior solution is (85), under $0 < T_k < D_k$ and $a_k T_k < D_k$. At origins with no illegal extraction the action set is empty, so the solution is $\pi = 0$ regardless of a hypothetical financial spread.

Integrating actual incremental payoffs over the cutoff gives

$$O_k = hX_k \int_0^{c_k \pi_k} \frac{c_k \pi_k - z}{c_k} dz = \frac{hc_k \pi_k^2 X_k}{2}, \quad (159)$$

$$F_k^{\text{real}} = hX_k \int_0^{c_k \pi_k} \frac{z}{c_k} dz = \frac{hc_k \pi_k^2 X_k}{2}. \quad (160)$$

No organizer surplus is omitted, nor assigned automatically to mineral research owners or local governments. Common ownership would require consolidating the accounts and then solving that owner's changed decision.

D.3 Global price solution and regularity

For safe origins $\pi_k = R_k = 0$; for the fixed smooth vulnerable branches $\pi'_k = -1/D_k < 0$. Since $l_{Ek} \geq 0$ and $e_k < 1$,

$$R_{k,\pi} = r_k l_{Lk} - A_k e_k > 0 \quad (161)$$

except the harmless zero-loss limit. Every term of F_{cw} and H_{cw} is nonincreasing in x , and

$$\frac{d}{d\pi} \frac{1 + \phi + R_k(\pi)}{1 - \pi} = \frac{R_{k,\pi}(1 - \pi) + 1 + \phi + R_k}{(1 - \pi)^2} > 0. \quad (162)$$

Thus $T' \leq 0$. Put $Y_0(x) = B_D x^\eta - H_{cw}(x) > 0$. Its derivative is $\eta B_D x^{\eta-1} - H'_{cw} > 0$, and

$$J'(x) = \frac{Y_0 + xY'_0}{T} - \frac{xY_0 T'}{T^2} > 0. \quad (163)$$

Theorem 6.1 now follows without assuming a contraction of buyers', fraudsters' and suppliers' separate best responses. The selected interval carries all origin-specific conditions $B_{Dk}x^\eta > B_{Ck}\pi_k/(1 - \pi_k)$. A compact primitive family with positive lower bounds on x , T , residual delivery, capacities and all strict margins yields a positive lower bound on J' . All formula derivatives are continuous there, giving bounded price and value derivatives by implicit differentiation.

Existence is asserted for that explicit branch. With new entry, an exhausted dirty source, effort saturation or a change in ownership, the clipped fraud fixed point remains valid but the displayed smooth clearing problem may change. In particular one cannot keep creating origin slots after its actual illegal stock constraint binds. Rationing or alternative contract groups must then be priced and solved separately.

D.4 An exact equilibrium and arbitrarily small positive leakage

At the parameters of Example 6.2, $a = 1$, $D_k = 5/6$, $\pi(x) = 1 - 6x/5$ and $x^* = 2/3$. The interval $[13/20, 7/10]$ lies strictly inside the effort and fraud branches and has positive origin-specific residual deliveries. At its endpoints the strictly increasing J straddles 40/139. Then $R = 14/75$, $T = (4/15)(139/75)/(4/5) = 139/225$, and $h^2 = (139/25)/(139/225) = 9$. Substitution also recovers the specified bonds and losses, establishing a feasible equilibrium with these parameters.

For any desired $\epsilon \in (0, 1)$ fix a compact positive price-ratio interval below $1 - rb$, keep reporting parameters and b fixed, and choose c large enough that

$$0 < \pi(x) = \frac{1 - x - rb}{c + b(1 - r)a} \leq \epsilon \quad (164)$$

throughout it. Choose B_D large enough for the strict origin residual conditions and choose E_D/E_C strictly between the resulting endpoint values of J . The class remains nonempty and interior; $e > 0$ and all obligations have actual separate funding. For a current cohort,

$\iota = \pi(1-p)/(1-\pi p) \leq \pi/(1-\pi) \leq \epsilon/(1-\epsilon)$. Taking the desired bound in this inequality constructs any arbitrarily low strictly positive live contamination level. This is a costly-forgery, funded-credibility restriction, not a claim that increasing a nominal fine with no wealth achieves low leakage.

D.5 Four linked accounts

For every origin and vintage keep four distinct state vectors. The material vector records physical quantities and current custodians, including material embedded in final products. Transfers and transformations preserve extraction origin and the material balance; an admissible transformation partitions parents into weighted descendants without making two full claims on the same input. The status vector identifies the currently live claims. An effective decision changes this vector for all affected descendants and their surviving physical weights. The financial vector contains the commodity escrow, distinct stake, promised compensation, claims horizon, and each funding source. The evidence vector retains the authenticated ancestry and decisions permanently.

At the current-date destination split, the physical identity is

$$Q_{Ck} + Q_{Dk} = X_k(1 - \pi_k p_k) + X_k \pi_k p_k + Y_k. \quad (165)$$

The first term is live certified material; its illegal part is $X_k \pi_k (1 - p_k)$. The second is excluded material, not negative extraction. Older holdings have their own v and cannot acquire the current government's legal status retroactively. Sum expectations of surviving masses over exactly the same origin-vintage-date set to obtain (92). Neither the expectation of individual ratios nor a current production weight generally equals this statistic. A zero denominator leaves the conditional rate undefined.

The slot commodity cash identity, including all fraudulent realizations in expectation, is

$$hX_k = h(1 - \pi_k)X_k + p_D \pi_k X_k + h b_k p_k \pi_k X_k + F_k^{\text{real}} + O_k. \quad (166)$$

Indeed the leftover spread is $h\pi_k X_k(1 - x - b_k p_k) = h c_k \pi_k^2 X_k$ by the cutoff equation. The sanction term is the organizer's expected forfeiture plus compensating bribes that finance the signers' expected forfeitures. Bribe receipts and signer bond losses cancel in the signers' net cash accounts; the public collector receives the actual expected forfeitures once. Minimum bribe payments use separate advance finance; expected solvency is not substituted for that liquidity requirement.

Across dirty suppliers the two commodity receipt channels sum to

$$p_D \sum_k Y_k + p_D \sum_k \pi_k X_k = p_D \sum_k Q_{Dk}. \quad (167)$$

Subtracting their private implementation cost gives (91). The clean suppliers' receipts similarly are $h \sum_k Q_{Ck}$. The fee budget is separate:

$$h\phi X = \text{paid ordinary service and capital bills} + \text{retained funded service balance.} \quad (168)$$

The buyer's total generalized bill is $h(1+\phi)X + h \sum_k X_k R_k = E_C$; only its first term is the invoice. Real forgery, diligence, incident handling, ordinary verification and production inputs enter resource use once. Taxes, fines, bribes, rents, fees and principal transfers do not create additional real resources. Household welfare includes their actual wages, fixed-factor rents, consumption, effort and environmental consequences. It is distinct from the government's fiscal value and both operators' net receipts.

D.6 Scaling, owners and benchmark recovery

If every regional capacity coefficient and both demand coefficients (by replicating the buyers) are multiplied by $m > 0$, the price ratio, prices and per-unit decisions are unchanged; output and every aggregate cash and workload entry multiply by m . This is replication of independent populations, not a representative auction. If coefficients multiply by $L^{1+\eta}$ while expenditure stays fixed, x is unchanged, prices divide by L , physical flows multiply by L , and aggregate invoice flows are unchanged. The invoice-indexed losses and real-cost assumption then require the explicitly paid productivity path; fixed technology at positive marginal physical verification expense cannot implement this scaling.

An owner's private success changes only its own dated capacity and liabilities, evaluated at the actual aggregate prices. In this common branch its operating value is linear in its own β_r at fixed prices, with coefficient $h^{1+\eta}/(1+\eta)$ or $(xh)^{1+\eta}/(1+\eta)$. Different contract wedges give different coefficients. The simplification follows from individual optimization; it is not inferred from a world supply index. The nonlinear proof uses its own actual optimized value instead.

Setting all vulnerable action sets empty gives $\pi = e = R = 0$, $X = B_C h^\eta$, $Y = B_D p_D^\eta$, clean generalized price $(1+\phi)h$, and both independent public-price power-supply markets. Replacing the quasilinear log expenditure system by the previously stated CES demands recovers the old CES benchmark only when its exact preferences, fee incidence, zero-leakage routes and continuation-linear assumptions are restored. No CES derivative is asserted for the present changed standardized contract without deriving that demand again.

E Proofs for the active nonlinear world

E.1 The cell, its regional cost, and its true physical account

All markets below are atomless replications of cells. A Muir cell itself has the continuum unit-demand buyers of the retained mechanism theorem. A VCG cell has exactly two whole buyers, each with capacity one. Replication does not split a whole buyer or take a continuum limit of an auction. There are exact aggregate frequencies of independent within-cell value profiles and resale events. Cross-cell reports and trade are unavailable; cell outcomes have no aggregate price effect. Conditional on the public state, the finite cell's two private values have the printed product law. These information and exact-aggregation assumptions are part of the institution; aggregate demand shocks would require a contingent-price game.

Producer access is also a primitive of this class. Before participation and procurement, the retail origins are geographically restricted to the assigned processing network, or face an enforceable origin-specific route restriction with the custody control specified below. They cannot legitimately deliver raw mineral to the bulk outlet, including by declining the retail consignment contract. Their participation alternative is exit at zero; the only positive-stock routes are the actual cell mechanism and the genuine retained conversion. Covered unauthorized route changes require the custody theorem's knowingly participating bonded veto parties and are deterred there; no unilateral raw-market bypass is admitted. The common final-product demand system does not expand this raw-access set. Without this restriction, the bulk outlet must enter W_r and can dominate the displayed retail optimum; separate group labels alone would not justify excluding it.

For a region with cell measure n_r , implementation costs add. If a physical allocation $z(a)$ to its cells produces Q_r , Jensen's inequality gives

$$\int \frac{k}{2} z(a)^2 da \geq \frac{k}{2n_r} Q_r^2, \quad \int z(a) da = Q_r, \quad (169)$$

with equality at $z(a) = Q_r/n_r$. More generally the power-cell cost printed in the main text gives the cost coefficient $\beta_r = n_r k^{-\eta}$. The public regional consignment contract requires equal stock per installed cell at every chosen total quantity, so its actual continuation is $W_r(Q_r; p) = n_r W_j(Q_r/n_r; p)$. Consequently

$$\max_{Q_r} \{W_r(Q_r; p) - kQ_r^2/(2n_r)\} = n_r \max_z \{W_j(z; p) - kz^2/2\} = n_r \pi_j(p).$$

This is exactly the original regional $W - \mathcal{C}$ interface with an attained cost fiber. It keeps the fixed retail population and routing rule in the inherited state. For comparison, even if independent cell stocks and their corresponding separate household technologies were allowed, each cell's profit is bounded by its attained one-cell maximum. Integrating gives the same upper bound, and identical maximizing choices attain it. Neither argument replaces a freely reallocated regional continuation by a separately optimized aggregate

cost envelope: without the routing rule or the separate cell costs those two extrema need not be jointly attainable. No representative auction is being inferred.

For the explicit quadratic household implementation, let each cell's household produce $q = a_0\sqrt{g}$, with $a_0 = 2/\sqrt{k}$, input resource price one, pass-through necessities and independently secured subsistence. The least-cost gate quotes are $p_g = 1$ and $\theta_B = kq/2$; input is $g = kq^2/4$. Indeed household optimization gives $\theta_B q = 2p_g g$, while the gate's bill net of input margin is $(p_g + 1)g$, minimized over $p_g \geq 1$ at one. Its cost is $kq^2/2$ and household operating rent is $kq^2/4$. The construction is measurable in (k, q) and integrates over finite cell and region measures because the finite public types and compact strict price branches bound k , controls and the resulting household inputs uniformly. It implements the installed-equipment branch of Γ_r ; endogenous entry, distress and labor participation outside that branch require their own feasible fibers. Gate payments are private implementation cost, not the real resource cost g or household welfare.

At a cell, primary physical mineral is m and retention is y , so $q = m + y$. A secondary sale merely reallocates part of m . If processing productivity is v , final output from primary/resale holdings is the integral or sum of v over final users. The retained route adds $H(y)$. One may have many final-product units per mineral unit: conversion factors are explicit, so final-product output is not itself a mineral mass. All mineral is embedded in its final product or retained-route product; there is no disposal, recycling or extra mineral from resale in this class. If a retained technology has saturated, its excess mineral remains recorded as terminal inventory; the examples' optimizing routes do not have such excess. Product material keeps its origin and certification history.

E.2 A strictly active Muir cell with positive retained delivery

Set $\rho = 1/50$, regulated mass $r = 1/200$, regulated productivity $M = 2$, and repeated acquisition charge $p\tau_0$, with $\tau_0 = 1/20$. The unregulated mass is one with inverse productivity demand

$$P(u) = \begin{cases} 1 - 8u/3, & 0 \leq u \leq 1/4, \\ 4(1 - u)/9, & 1/4 \leq u \leq 1. \end{cases} \quad (170)$$

Write

$$a = \frac{6 + \sqrt{6}}{48}, \quad b = \frac{1 + \sqrt{6}}{8}, \quad s = 1 - \frac{16a}{3}, \quad c = \frac{8}{3}a^2, \quad R(u) = uP(u). \quad (171)$$

The least concave majorant is $\bar{R}(u) = c + su$ on $[a, b]$ and $R(u)$ elsewhere. The two quadratic pieces have the same tangent at a, b ; they lie below it, and the joined slopes are nonincreasing. This proves the majorant, including minimality at its contact points.

For completeness the exact fixed-volume receipts are

$$V(m) = \begin{cases} (M - \tau_0)m, & 0 \leq m \leq r, \\ (1 - \rho)[\bar{R}(m - r) + r(M - \tau_0)] + \rho m P(m - r), & r < m \leq 1 + r. \end{cases} \quad (172)$$

To establish this in the actual game, for $m > r$ any secondary price above one generates unregulated supply exceeding all possible regulated demand by at least $m - r$. Since $M - \tau_0 > 1$, all regulated buyers eventually own, and the untaxed secondary market clears uniquely at $P(m - r)$. An unregulated primary unit then has incremental value $(1 - \rho)v + \rho P(m - r)$; a regulated unit has incremental value $(1 - \rho)(M - \tau_0) + \rho P(m - r)$. If regulated primary service is d , the threshold-envelope upper bound on receipts is $(1 - \rho)[\bar{R}(m - d) + (M - \tau_0)d] + \rho m P(m - r)$. Its slope in d is strictly positive because all slopes of $\bar{R}$ are at most one. Thus $d = r$; the two contact thresholds attain the bound and clear the same secondary market. For $m \leq r$, ownership values are at most $M - \tau_0$, attained by the regulated lottery and clearing reference price M . This proves the first line, including the downward right-hand jump possible at $m = r$. All participation is relative to actual resale access. This argument is the separated-group specialization of the retained continuum theorem, not a volume-only resale formula asserted for general regulated distributions.

Introduce

$$d = \frac{16\rho}{3}, \quad \alpha = (1 - \rho)s + \rho(1 + 8r/3), \quad m_0 = 41/200, \quad k = \frac{\alpha - dm_0}{m_0}, \quad (173)$$

and the independent, genuine retained conversion

$$H_\delta(y) = \begin{cases} h_0 y - y^2/(2\delta), & 0 \leq y \leq \delta h_0, \\ \delta h_0^2/2, & y \geq \delta h_0, \end{cases} \quad \delta = 1/1000, \quad h_0 = 1/5. \quad (174)$$

It cannot change a retail buyer's mineral outside option. Its final product is sold to the common final consumer, so its value is exactly pH_δ .

Lemma E.1 (Global Muir solution on an explicit price interval). *For every $p \in I = [99/100, 101/100]$, the global maximum in (93), with V as in (172), is unique and equals*

$$\begin{aligned} m(p) &= \frac{p\alpha - k\delta(h_0 - \alpha)}{k(1 + \delta d) + pd}, \\ y(p) &= \delta[h_0 - \alpha + dm(p)], \quad q(p) = m(p) + y(p), \\ \pi(p) &= p[V(m(p)) + H_\delta(y(p))] - kq(p)^2/2. \end{aligned} \quad (175)$$

It has $a < m(p) - r < 1/4$, $0 < y(p) < \delta h_0$, strictly positive profit, an active rationing premium and positive resale.

Proof. First put $H = 0$. Positive retention is then dominated by lower production. For $m > r$, the profit is piecewise quadratic. On $a + r \leq m \leq 1/4 + r$ its derivative is $p\alpha - (k + pd)m$, with unique zero $m_p^0 = p\alpha/(k + pd)$. On $r < m < a + r$, the objective is strictly concave and its joining derivative at $a + r$ is positive. On $1/4 + r \leq m \leq b + r$, its derivative is decreasing and at its left endpoint equals

$$p\{(1 - \rho)s + \rho(4/9)(1/2 - r)\} - k(1/4 + r) < -.007.$$

It therefore decreases throughout; the last piece is strictly concave with the same negative derivative at $b + r$. For certified elementary bounds, $2.44948 < \sqrt{6} < 2.44950$ gives $0.2845 < k < 0.2848$, $.2034 < m_p^0 < .2066$ and the displayed strict derivative bound uniformly on I . This covers every positive-volume branch and kink. Quantities $m \leq r$ earn at most $1.01(M - \tau_0)r < .01$, whereas evaluating at m_0 gives profit above .09. Thus the interior maximizer is globally unique.

The loss at the left endpoint of the winning interval is $(k + pd)(m_p^0 - a - r)^2/2 > 9 \cdot 10^{-5}$; the loss at the right endpoint exceeds $3 \cdot 10^{-4}$. The preceding monotonicities show that these bound losses against all $m > r$ outside the interval, and the small-volume gap is larger. The same rational bounds certify these inequalities.

Now restore H_δ . For every m, y the extra retained receipts are at most $p\delta h_0^2/2 \leq 2.02 \cdot 10^{-5}$, while producing y adds a nonnegative cost. The feasible choice $y = 0$ preserves the old maximal value. Hence no new global optimum can leave the winning interval in m . For each such m the unique optimizing retention is $y = (ph_0 - km)/(p/\delta + k) > 0$, strictly below δh_0 . The joint objective on this rectangle has negative definite Hessian: the negative of its Hessian is

$$\begin{pmatrix} pd + k & k \\ k & p/\delta + k \end{pmatrix},$$

with positive determinant and diagonal. Its two first-order equations are $p(\alpha - dm) = k(m + y) = p(h_0 - y/\delta)$. Solving gives (175); the printed constants put this solution strictly inside the rectangle. It is therefore the unique global solution.

Put $u = m - r$ and $\lambda = (u - a)/(b - a)$. The unregulated menu serves the top a certainly and the next $b - a$ with probability $\lambda \in (0, 1)$; its payment transform is $z \mapsto p[(1 - \rho)z + \rho P(u)]$. Its premium over the same-volume efficient primary sale is $p(1 - \rho)[\bar{R}(u) - uP(u)] > 0$. Conditional resale volume is $\lambda(b - u) = (1 - \lambda)(u - a) > 0$. Every regulated buyer is already an owner, so no regulated secondary purchase occurs on this path, although such purchases remain feasible and taxed in the game. Retention is strictly positive by the preceding solution. Profit remains above .09. $\square$

Let $I_P(u) = \int_0^u P(v) dv$. Its explicit formula is $u - 4u^2/3$ for $u \leq 1/4$, and $I_P(1/4) + (4/9)[u - u^2/2 - 7/32]$ above $1/4$. The actual final output, with all conversion factors present, is

$$Z_M(p) = rM + (1 - \rho)\{I_P(a) + \lambda[I_P(b) - I_P(a)]\} + \rho I_P(u) + H_\delta(y(p)). \quad (176)$$

Without resale the lottery determines who processes mineral; with resale the top u unregulated users process it. This proves the expression and shows why replacing Z_M by q , V or $\bar{R}$ would be wrong. Direct differentiation gives

$$\begin{aligned} m'(p) &= \frac{k(\alpha + \delta dh_0)}{[k(1 + \delta d) + pd]^2} > 0, \\ q'(p) &= (1 + \delta d)m'(p) > 0, \\ Z'_M(p) &= m'(p) \left[(1 - \rho) \frac{I_P(b) - I_P(a)}{b - a} + \rho P(u) + \delta d H'_\delta(y) \right] > 0, \\ \pi'(p) &= V(m(p)) + H_\delta(y(p)) > 0. \end{aligned} \tag{177}$$

The last identity is the valid interior envelope, not a shortcut from physical quantity or auction exposure. Its revenue derivative differs from final output because recipients of retail rents and charges differ. Strict inequalities, global comparison gaps and Hessian nonsingularity survive small finite-dimensional perturbations of the displayed piecewise twice continuously differentiable primitives that preserve their contact and branch structure, including small paid institutional charges. This is not a claim about arbitrary uniformly small nonsmooth perturbations. Thus the active class has nonempty interior in that stated parameter class; the numerical example is one member of this class.

E.3 A separate finite whole-buyer VCG construction

One cell has two buyers with independent raw unit productivities $v_1, v_2 \sim U[0, 1]$, one indivisible package per buyer and package mass one. Direct values are pv_i , genuine seller retention is phq , with $h = 1/10$, and production costs $k_V q^2/2$, with $k_V = 1/5$. At every produced stock q , the seller uses the all-offered whole-buyer pivot auction with $Q = \lfloor q \rfloor$ available packages and the complete terminal consent protocol. All buyers, including nonparticipants, can join terminal trade. The genuine public retention is tradable through that protocol, all consent protects initial payoffs, and equilibrium terminal allocations are efficient. Truthful reports and no terminal trade attain the retained efficient-complete-game bound. This uses the original independent zero-inclusive buyer domain and whole-buyer removal; it is not a monopoly theorem outside that comparison class.

Lemma E.2 (Nonlinear finite continuation and actual production). *Put $R_h = h + (1 - h)^3/3$. The actual all-offered continuation is*

$$W^V(q; p) = \begin{cases} phq, & 0 \leq q < 1, \\ p[h(q - 1) + R_h], & 1 \leq q < 2, \\ phq, & q \geq 2. \end{cases} \tag{178}$$

For every $p \in I$, production is uniquely $q_V = 1$, actual net profit is $\pi_V(p) = pR_h - k_V/2$, expected retained mineral is $h^2 > 0$, and actual final output is $Z_V = (2 + h^3)/3$.

Proof. With no available whole package everything is retained. With one package, the highest buyer receives it if and only if its productivity exceeds h , and its whole-buyer pivot payment is $p \max\{h, \min(v_1, v_2)\}$. If neither buyer exceeds h , the whole unit is retained instead. Adding true retention of $q - 1$ units makes total expected value

$$p \left[h(q - 1) + h + \int_h^1 \Pr(\min(v_1, v_2) > t) dt \right] = p[h(q - 1) + R_h].$$

With at least two available packages each buyer above h pays exactly ph , and every other unit is retained. Thus total seller value is phq . The change at $q = 2$ is a real effect of the committed all-offered inventory, not a reason to replace W^V by a smooth envelope or a one-buyer pivot.

On the retention-only branches, $phq - k_V q^2/2$ is bounded above by $(ph)^2/(2k_V) < .026$. On $1 \leq q < 2$, its derivative is $ph - k_V q < 0$, so the maximum is at $q = 1$. There its value is $p(.343) - .1 > .239$, proving the strict global optimum. The stock is retained with probability $\Pr(\max(v_1, v_2) < h) = h^2$; otherwise one buyer processes it. Expected final output equals

$$\mathbb{E} \max\{h, v_1, v_2\} = h + \int_h^1 (1 - t^2) dt = (2 + h^3)/3.$$

It counts the retained conversion when no sale occurs. Since the initial allocation is efficient, no mutually beneficial terminal reallocation exists on the truthful path. Resale remains accessible; active equilibrium trade is not inferred from accessibility. The complete-game attainment conditions stated before the lemma preserve this continuation against combined unilateral initial and terminal deviations. $\square$

This finite example is nonlinear in stock, although its selected strict branch is affine in the common final-product price. It is deliberately distinct from the active resale and inefficient primary lottery of Muir. Whole-package capacity also explains why general differentiable supply elasticities cannot simply be asserted for every finite auction.

E.4 The equilibrium operator, reconstruction and scale

Theorem E.3 (Normalized retail equilibrium and actual regional values). *Let finitely many types have the preceding strict cell branches, or sufficiently small perturbations preserving their global optimality and $Z'_j(p) \geq 0$. Let (n, E_R) range in a compact set with $n_j \geq 0$, $\sum_j n_j \geq \underline{n} > 0$, $E_R > 0$, and*

$$p_- \sum_j n_j Z_j(p_-) < E_R < p_+ \sum_j n_j Z_j(p_+), \quad [p_-, p_+] \subseteq I. \quad (179)$$

There is a unique clearing price in this interval; the selected equilibrium price, quantities, net profits and measurable household allocations depend continuously on the state. The price is continuously differentiable where the parameterization is differentiable, including

a smooth extension to zero mass of any individual type. Regional and type aggregates are exactly

$$Q_j = n_j q_j(p), \quad \Pi_j = n_j \pi_j(p), \quad Z_R = \sum_j n_j Z_j(p). \quad (180)$$

Scaling all region cell measures and E_R by $L > 0$ leaves price and cell choices unchanged and multiplies these aggregates by L .

Proof. The left side of clearing is continuous and has derivative $\sum_j n_j [Z_j(p) + pZ'_j(p)] > 0$, uniformly bounded below on the compact domain. Endpoint signs give a root by the intermediate value theorem and strict monotonicity gives uniqueness in the stated interval. Its derivative is nonzero, so the implicit function theorem gives local differentiability; uniqueness identifies the local solutions and compactness gives a global continuous selected branch. Zero mass of a type creates no singularity because the remaining positive core keeps the derivative positive. The formulas for counterfactual individual cells remain defined there.

Each cell's global production and mechanism optimum and the consumer's strictly concave demand optimum are attained. The atomless price-taking assumption makes these choices individually optimal at the common root. Clearing exhausts final production. Every region attains the integrated cell bound already proved, and the constructive household map proves reconstruction. Bounded prices, cell controls and finite masses give integrability; finite public types give measurable selection. Thus these choices form the specified economy's equilibrium. Finally dividing its clearing equation by L proves homogeneity and the value formulas. $\square$

The endpoint domain is nonempty: choose any positive cell masses and $p_0 = 1$, set $E_R = p_0 \sum_j n_j Z_j(p_0)$, and take an interior interval around p_0 . Strict monotonicity gives both strict signs, which persist under small heterogeneous perturbations. For a vanishing dirty retail population retain any positive clean cell core. The individual physical response is weakly increasing in p for VCG and strictly increasing for Muir. Collective reform or innovation affects the price through the *whole* final-output sum. For fixed E_R , more cells lower price strictly; cell output and protection profit then move by their actual functions. No derivative sign follows if demand also changes arbitrarily.

On paths $n_{j,t} = L_t \hat{n}_{j,t}$, $E_{Rt} = L_t e_{Rt}$ inside this compact domain, normalized profits and physical controls are uniformly bounded. If $L_{t+1}/L_t \leq g$ and $\delta_* g < 1$, any owner's bounded per-unit-scale discounted continuation has a uniform tail bound $C(\delta_* g)^T / (1 - \delta_* g)$. The owner mass does not become the adopter population. Replicating an independently licensed and financed cell gives an exact own value increment; public copying cancels later terms only if the same liabilities and financing remain after both histories. This is the nonlinear stage interface used by the dynamic proof.

E.5 Proof of the explicit certification connection

Use disjoint finite origin/contract groups for bulk supply and retail cells, all of which have their regional cost fibers and physical origin records. Let θ collect their normalized capacities and all publicly stated contract parameters. At $\varepsilon = 0$, let the bulk subsystem's two price equations have a continuously differentiable solution $(h_0(\theta), x_0(\theta))$, with nonzero Jacobian determinant on a compact neighborhood. Its solved common-world theorem supplies this property away from clipping switches and boundary deliveries. Retail clearing supplies $p_0(\theta)$ and its strictly positive derivative by Theorem E.3. Require common positive margins for all procurement, service, exposure, reporting, finance and physical bounds; fees, reserve returns and fraud-organizer income are the actual accounts from that subsystem, not omitted transfers.

Define residuals

$$\begin{aligned} R_C &= hK(x)X - E_C, \\ R_D &= hxY - E_D - \varepsilon Y / (Y + \gamma Z_R(p)), \\ R_R &= pZ_R(p) - E_R - \varepsilon \gamma Z_R(p) / (Y + \gamma Z_R(p)). \end{aligned} \tag{181}$$

Here X, Y are the *same* endogenous quantity functions as in the bulk certification theorem; Z_R is (176) or its finite VCG counterpart, summed over actual cell masses.

Proof of Theorem 7.2. The utility in (95) is strictly concave on the positive orthant: the three log terms with positive coefficients already have a negative definite Hessian and the last log is concave. Its marginal willingness to pay is

$$P_X = E_C/X, \quad P_Y = E_D/Y + \varepsilon/(Y + \gamma Z_R), \quad P_R = E_R/Z_R + \varepsilon\gamma/(Y + \gamma Z_R).$$

Consequently the first-order equations are necessary and sufficient for the final consumer's global optimum. Effective certified willingness $P_X = hK(x)$ includes the original composition-dependent losses and institutional charges. Substituting $P_Y = hx$ and $P_R = p$ gives exactly (181). Total final expenditure is $E_C + E_D + E_R + \varepsilon$ by multiplication of each marginal value by its quantity and summation; sufficient consumer numeraire funds it. Retail primary and resale cash transfers redistribute this expenditure and do not create another demand budget. Procurement transfers go to the reconstructed households, service payments to their specified payers, and charges to the stated authority. Binding escrow and exposure budgets remain those of the supplied contract game.

At $\varepsilon = 0$ the derivative of (R_C, R_D, R_R) with respect to (h, x, p) is block diagonal. Its first block is the assumed nonsingular bulk Jacobian, and its last entry is $\sum_j n_j [Z_j + pZ'_j] > 0$. Thus its determinant is nonzero. All residuals are continuously differentiable on the stated strict domain. The implicit function theorem gives a unique differentiable equilibrium branch near each baseline point for small ε . Restricting to a sufficiently small closed state neighborhood inside one such chart gives a common positive upper bound for ε and compact continuous branch. Equivalently a finite covering gives a common local result on

a compact base branch when its neighborhoods and uniqueness regions are compatible. No global uniqueness of all possible regimes is needed or claimed.

Strict cell optima, menu premium, resale volume, retained delivery, funded reporting and procurement, and positive X, Y, Z_R survive by continuity. The bulk actions remain their actual best responses at the new (h, x) because their contract equations were not replaced. The retail menus and quantities remain their actual best responses at the new p . The consumer's optimum has already been proved and every clearing and physical equation holds. These choices therefore form the complete connected equilibrium. The connection is nonzero: at $\varepsilon > 0$, $\partial P_Y / \partial Z_R = -\varepsilon \gamma / (Y + \gamma Z_R)^2 < 0$ and $\partial P_R / \partial Y$ has the same strictly negative value. It cannot be removed while preserving the stated final demand. $\square$

This proof is a nonempty extension, since the bulk class has strict interior examples, Lemmas E.1 and E.2 give strict retail examples, and their product satisfies the block-diagonal hypotheses. Taking a sufficiently small positive ε yields genuine cross-market substitution and retains every strict mechanism effect. Heterogeneous positive cell scales and a small positive reformable population can be selected in that open neighborhood. Stronger global claims require more than this local result.

If all capacities, all expenditures and ε scale by L , division of (181) by L gives the same normalized system; its regularity and uniform bounds persist while that state remains in the verified compact domain. If $Y/L \rightarrow 0$ while E_D/L stays bounded below, the equation $hx \geq E_D/Y$ forces the dirty price to diverge. Thus that path does *not* satisfy this theorem's compact-price hypotheses. A boundary theorem must instead scale outside spending appropriately and verify its rescaled Jacobian, or preserve positive outside demand and supply cores. Discounting alone does not repair a missing stage equilibrium bound.

E.6 A uniform connected stage for an infinite growth path

The preceding boundary warning does not prevent growth. The following mixed-scale class solves the boundary with an explicit positive outside core and fixed vulnerable origins. Use a growth scale $L \geq 1$ and put $\lambda = L^{-1} \in (0, 1]$. Safe certified production has coefficient Lb_s , where $b_s > 0$; a fixed vulnerable origin group has coefficient $\beta_v > 0$; outside dirty production has coefficient $B_D > 0$ in a compact positive interval. The latter groups retain their absolute physical masses. They are not interpreted as disappearing merely because division by L makes their normalized mass vanish. Retail cell masses are $L\hat{n}_j$, with a positive clean retail core. Write $\hat{n}_j = n_j/L$ and $\hat{n} = (\hat{n}_j)_j$ for their normalized masses. Their normalized final output is $Z(p; \hat{n}) = \sum_j \hat{n}_j Z_j(p)$.

Set certified expenditure $E_C = Le_c$, retail expenditure $E_R = Le_r$, outside expenditure $E_D = e_d > 0$, and keep the utility coefficient $\varepsilon_u > 0$ fixed in (95). The safe contracts have

no illegal-origin admission; the vulnerable contract has the common-world game's actual interior fraud share $\pi(x) \in (0, \bar{\pi})$, loss/loading term $\mathcal{R}_v(x)$, and funded fee $\phi > 0$. Write

$$\begin{aligned} t_v(x) &= \frac{1 + \phi + \mathcal{R}_v(x)}{1 - \pi(x)}, \\ X &= h^\eta \left[Lb_s + \frac{\beta_v}{1 - \pi(x)} \right], \\ Y &= h^\eta \left[B_D x^\eta - \frac{\beta_v \pi(x)}{1 - \pi(x)} \right]. \end{aligned} \tag{182}$$

All reporting, finance and exposure parameters retain their paid, invoice-scaled interpretation. In particular β_v is actual origin capacity, not a probability weight substituted for material. Every retail input group and every bulk input group is disjoint in the origin/contract partition. If one jurisdiction owns both branches, its installed capacity is explicitly split between them and its costs and receipts are added once.

Theorem E.4 (Uniform mixed-scale connected operator). *Choose a compact state rectangle $s = (b_s, B_D, \hat{n}, e_c, e_d, e_r)$ near a positive baseline, with a positive retail core, retail endpoint margins, and a compact interval of baseline x strictly inside the fraud branch. Suppose the underlying reporting, funding and origin-stock margins are strict there. For all sufficiently small fixed $\beta_v > 0$ and $\varepsilon_u > 0$, the normalized equations*

$$\begin{aligned} 0 &= h^{1+\eta}[(1 + \phi)b_s + \lambda\beta_v t_v(x)] - e_c, \\ 0 &= xhY - e_d - \frac{\varepsilon_u \lambda Y}{\lambda Y + \gamma Z(p; \hat{n})}, \\ 0 &= pZ(p; \hat{n}) - e_r - \frac{\varepsilon_u \lambda \gamma Z(p; \hat{n})}{\lambda Y + \gamma Z(p; \hat{n})} \end{aligned} \tag{183}$$

have a unique continuously differentiable selected branch $(h, x, p) = \Phi_{\beta_v, \varepsilon_u}(s, \lambda)$ in a common neighborhood of the baseline solutions, uniformly for $\lambda \in [0, 1]$. Prices, absolute Y , Z , active Muir premium, resale, retained deliveries, funding and origin-stock margins have uniform positive bounds on this compact branch. Every finite L has positive endogenous vulnerable admission and nonzero cross-market substitution. The equations retain their exact absolute physical accounts at $\lambda = 0$ as limits.

On a slightly smaller compact rectangle, if the primitive functions are twice continuously differentiable on their strict branches, then

$$\sup \|D\Phi_{\beta_v, \varepsilon_u} - D\Phi_{0,0}\| \leq C(\beta_v + \varepsilon_u). \tag{184}$$

For the retail price alone, with $p^0(\hat{n}, e_r)$ the scalar retail root, the stronger bound $\sup \|D(p - p^0)\| \leq C\varepsilon_u$ holds. Consequently its error in a one-period price change is bounded by $C\varepsilon_u \|(s', \lambda') - (s, \lambda)\|$. It is not a fixed error independent of the size of the state change.

Proof. Equation (182) is obtained by summing actual regional safe supply, vulnerable slots and residual dirty mineral. Multiply the three consumer marginal conditions used in the preceding proof by their quantities, substitute $E_C = Le_c$, $E_R = Le_r$, $Z_R = LZ$, and divide the certified and retail equations by L . The dirty equation is kept in absolute units. This yields exactly (183). Its denominator stays positive at $\lambda = 0$ because Z has a positive core.

At $\beta_v = \varepsilon_u = 0$ the explicit baseline is

$$h^0 = \left[\frac{e_c}{(1+\phi)b_s} \right]^{1/(1+\eta)}, \quad p_D^0 = \left[\frac{e_d}{B_D} \right]^{1/(1+\eta)}, \quad x^0 = p_D^0/h^0, \quad p^0 Z(p^0; \hat{n}) = e_r.$$

The first two residuals are then $(1+\phi)b_s h^{1+\eta} - e_c$ and $B_D(xh)^{1+\eta} - e_d$. Their derivative in (h, x) is lower triangular with strictly positive diagonal; the retail derivative is $Z + pZ' > 0$. All diagonal entries and the determinant are uniformly separated from zero on the compact baseline rectangle, independently of λ . Prices and outside mineral have uniform positive bounds.

Here is a uniform version of the implicit-function argument. Let $A(s)$ be this baseline price Jacobian at $\Phi_{0,0}(s)$; its inverse is uniformly bounded. On a common sufficiently small ball around each baseline price vector, uniform continuity of the baseline derivative makes $\|I - A(s)^{-1}D_u R_{0,0}(u; s, \lambda)\| \leq 1/4$. Compactness and smoothness of (183) then permit common positive bounds on β_v, ε_u such that the same quantity for the perturbed residual is at most $1/2$. Make those bounds smaller so the displacement $\|A(s)^{-1}R_{\beta_v, \varepsilon_u}(\Phi_{0,0}(s); s, \lambda)\|$ is at most half the ball radius. The map $u \mapsto u - A(s)^{-1}R_{\beta_v, \varepsilon_u}(u; s, \lambda)$ then maps that ball into itself and is a contraction with modulus $1/2$. It has a unique fixed point, uniformly for every (s, λ) in the compact domain. Ordinary differentiation of its nonsingular residual gives the continuous derivative; the common contraction and uniqueness identify all charts. Strict physical and institutional inequalities persist after reducing the perturbation bounds once more.

Nonemptiness is constructive. Choose any interior fraud ratio x_* from the common-price example, $h_* > 0$, positive b_s, B_D , and set $e_c = (1+\phi)b_s h_*^{1+\eta}$ and $e_d = B_D(x_* h_*)^{1+\eta}$. Choose positive retail masses, $p_* = 1$, and $e_r = Z(1; \hat{n})$. The required baseline equalities and strict retail margins hold; sufficiently small state neighborhoods preserve the ratio interval. At positive small β_v , actual vulnerable admission is $h^\eta \beta_v \pi(x)/(1 - \pi(x)) > 0$. It is uniformly smaller than the available absolute dirty output by the positive stock margin. The safe certified stream is $Lh^\eta b_s$, so contamination is computed with the correct increasing physical denominator. Reporting remains endogenous because x enters the actual fraud best response. Cross substitution is strictly nonzero whenever $L < \infty$ and $\varepsilon_u > 0$. At the limiting $\lambda = 0$, the dirty market still clears in absolute units against e_d and includes its fixed vulnerable diversion; it has not been discarded from the state or the accounts.

On a compact strict branch the residual and its first derivatives differ from their unperturbed values by $O(\beta_v + \varepsilon_u)$ uniformly. Uniformly bounded second derivatives, in-

verse Jacobians and the price displacement bound imply (184) by differentiating $D\Phi = -(D_u R)^{-1} D_s R$. For the retail component set $\varepsilon_u = 0$ while keeping the chosen β_v : its third residual is exactly scalar retail clearing, so its price is precisely p^0 , independent of $b_s, B_D, e_c, e_d, \lambda$. The same derivative argument with ε_u alone gives the stronger retail bound. Integrating the derivative of $p - p^0$ along the line segment between two states in the rectangle gives the final one-period comparison. This justifies decaying-drift estimates along $\lambda_t = g^{-t}$; a uniform level-error bound alone would not suffice. $\square$

The theorem preserves a fixed positive auxiliary outside-mineral core. If that core contains illegal extraction, it gives an explicit floor to world illegal output. A physical conservation target must include the floor, or identify these auxiliary origins as outside the protected ecosystem and assess the ecosystem-specific criterion. Neither zero normalized mass nor a falling contamination fraction proves zero absolute extraction. Safe bulk output and retail output can nevertheless grow at scale L , requiring the fully funded verification and permanent-record capacity established in Appendix F.

E.7 A financed faithful retail provenance subclass

The nonlinear retail cells also need physical provenance and paid service. The following subclass supplies them while the separate bulk vulnerable group retains strictly positive endogenous leakage. Every retail cell has a registered extraction site, dated physical cohort, producer/custodian, inspector and downstream processors. True extraction origin is immutable; governance changes affect new cohorts from the next production date and cannot relabel old dirty mineral. Physical stock measurement and a custody lock before buyer reports are assumed available and paid. Only measured finite stock enters the registered menu; outside mineral cannot enter after that lock. Alleged origin and collateral eligibility remain subject to forensic challenge. The actual procurement contract, retained route and verified pre-report stock contingencies are committed before measurement. No primary invoice or lottery obligation precedes branch selection. Thus feasible pre-lock diversion uses the original contractual branch; post-lock physical removal is infeasible in this subclass.

Proposition E.5 (Funded truthful retail admission). *On the compact stage domain of Theorem E.4, assume bounded retail and replacement/import invoices per registered cell and the preceding finite-stock deviation domain. There is a nonempty financed institutional class in which no coalition of producer, custodian, inspector and reachable downstream processors profits from false eligible-origin admission, finite diversion, selective suppression or their joint actions. Retail profit is its actual faithful cell profit. Bulk organizer profit remains in the separate bulk accounts.*

Proof. Specify evidence access. The custody system keeps redundant append-only measurement and origin records. The inspector and each processor can read its batch evidence; an independent forensic authority retains access after sale and after private silence. For every nontrivial admissible fraud affecting a cell, this authority obtains legally usable

contradictory physical evidence and collects the stated exposed assets with conditional probability at least $r_0 > 0$, after any selective suppression or mitigation choices, before the relevant funded claim expires or its collateral pool is released. This is an explicit unsuppressible-late-enforcement and physical access primitive, whose audit, archive, adjudication and collection costs are paid below. It is not an inference from ledger consensus. A private report can accelerate enforcement and earn a funded reward, but cannot turn the late audit off. Independence across different cells is unnecessary.

Let $\bar{q}$ bound cell extraction and $\bar{F}$ bound all coalition members' additional final-product sales, retained private benefits, feasible invoice receipts, avoided collectible fees and stolen receipts in one simultaneous cell action. For pre-procurement false-origin actions include any avoided implementation or procurement costs as well; for the post-procurement residual comparison sunk procurement still cancels. Registered stock limits, compact common prices, bounded costs and invoice obligations and limited retail access give a finite $\bar{F}$. Extra ledger units without measured backing are inadmissible. The Muir cell has at most $1 + r$ buyers, bounded productivity M , and retained receipts at most $p_+ \delta h_0^2/2$; VCG receipts and true retention are bounded by the finite stock and buyer values. Thus every originally committed residual branch has actual value in $[\underline{W}, \bar{W}]$. For the pre-report, no-invoice contract here take $\underline{W} = 0$; priced residual liabilities would instead remain in that bound. Formula (185), and its same-stock false-origin counterpart, give

$$G(a) \leq \bar{F} + \bar{W} - \underline{W} =: G_* < \infty.$$

Sunk procurement cancels and residual scarcity gains remain included. Secondary invoices within the processors' coalition cancel from its common gain and are not extra fraud resources.

Every false admission needs an indispensable independent licensed custody or inspection signer with a distinct collectible cell pool B ; it may be split between those service providers. Pool beneficial ownership belongs solely to these legally separate parties, never to the production operator or research owner, and is not transferable with a production license. The joint-control prerequisite is substantive: for every covered diversion, false-origin, invoice or suppression action, at least one such bonded custodian observes the true origin, physical movement and relevant action before authorizing it, has an effective physical/signature veto, and must knowingly participate for the action to succeed. Keys cannot be forged and the producer cannot bypass this custody authorization. This is satisfied by directly monitored registered origins and jointly controlled custody; it is not inferred for opaque imported provenance. An innocent inspector's unobserved exposure does not enter the deviating coalition's sanction. Each counted asset is legally exposed by the actual action and is counted once. For every covered action, the knowingly necessary participating signers' own collectively exposed assets must sum to at least B . Splitting a pool is allowed only if the necessary owners of its counted pieces must all join; a small participating share cannot use an innocent provider's balance. Choose $B > (G_* + 1)/r_0$. Expected collectible loss then exceeds the entire common gain, before nonnegative operating costs. Voluntary internally financed side payments cannot make all necessary participants better off. This

covers buyer bribes and selective suppression because the collection floor holds after those choices. For a joint action across cells take the union of their distinct pools and sum the gain bounds. Linearity of expectation suffices under correlated audits. A common principal may own several pools but cannot pledge one twice. Each remains segregated until its last claim closes; reuse requires replacement funding. These are asset constraints, not multiplied labels.

An actual payer is a national treasury with a separately allocated nonmineral numeraire endowment. It finances initial infrastructure and segregated pools, and pays honest service costs and a stated capital carrying payment. Pool assets are vested, legally collectible wealth of the independent signers while pledged; conviction destroys that signer's own vested claim. The treasury cannot reimburse forfeiture. Honest participation follows because service and carrying payments cover the stated outside options. Custody measurement, access, audits, storage, adjudication, maintenance and bounded report rewards all have positive finite prices. Report reserves are segregated beforehand; late enforcement has its own appropriation, so private silence cannot save its budget.

Let a_{init} bound infrastructure, vested pool grants and initial reserves per newly installed cell, and $a_{\text{run}} > 0$ bound service, capital carry, reserve replacement and current permanent-record storage per live cell, with the finite claims horizon stated in the dynamic contract. These constants are finite on the compact domain. The treasury must pay

$$\begin{aligned} \mathcal{A}_t \geq & a_{\text{init}}(\text{new funded cells at } t) + a_{\text{run}}(\text{live serviced cells at } t) \\ & + \text{past-record service and closed-cohort obligations.} \end{aligned}$$

Its nonmineral endowment funds this amount independently of hypothetical post-exit fee revenue. Earmarked service and capital transfers are not unrecorded additions to producer operating profit. Newly installed and live cell measures are bounded by constants times L_t on the growth class; the dynamic proof's paid verification productivity and permanent record account bound the complete remaining workload. Allocating the corresponding endowment and supplying initial reserves makes all budget inequalities strict. This constructs the promised nonempty financed class without changing the solved retail allocation or its net profit. $\square$

Retail false-origin organization has no profitable deviation because its actual finite-action payoff is negative; bulk vulnerable organization retains its separately solved positive surplus and low live contamination. Both sets of physical material, financial exposure and permanent evidence enter the consolidated accounts. The fixed auxiliary bulk extraction floor remains in every physical conservation comparison.

E.8 Commitment, limits and counterexamples

Let Σ record the original stock, menus, public information, allocations already promised, retained conversion, primary invoices, default/refund terms, and secondary market selec-

tion. At every compared history the stipulated continuation must have an equilibrium. Let Ω be the market state. The exact organizer's deviation comparison is

$$F_C(a) - k_f(a) + \mathcal{W}_\Sigma(q - d(a) \mid q, \Omega) - \mathcal{W}_\Sigma(q \mid q, \Omega), \quad (185)$$

before its action-specific sanctions and side transfers. Procurement is already paid and cancels. An originally committed menu for every verified pre-report stock, with no prior primary invoices and a custody lock after stock measurement, permits substitution of the corresponding optimized branch at that verified state. Hidden diversion after a lottery was sold does not. The outside conversion capacity and every liability remain in the inherited state. Bounded cell values on the compact domain provide finite oscillation bounds for the full coalition calculations; they do not by themselves fund the required collateral or service.

Three limits locate the result precisely. First, if every buyer can obtain the same mineral service at public quote r and retention delivers to that same outlet, voluntary primary receipts are at most rm and retention at most $r(q - m)$. Retaining everything attains rq , so the optimized value is rq . A strictly positive retail premium is impossible in that class. The access restriction and genuinely different processing opportunities above cannot be silently removed.

Second, not every nonlinear stage has a unique continuous equilibrium operator. One whole-buyer VCG cell may switch optimally between a retained stock and one offered package; at a tie, output jumps or is set-valued. The strict-gap interval in Lemma E.2 deliberately avoids this event. A prescribed discontinuous tie choice can fail to clear a market at the exact intervening demand. Atomless mixing among genuinely indifferent cells can restore convexified supply, but requires an explicit selection and is not a consequence of discounting.

Third, innovation and governance signs depend on the clearing institution. Even in the exact retail operator, reducing dirty cell mass with fixed clean mass and fixed E_R raises p , because the left side of (94) falls at every fixed price. Equations (177) then raise remaining Muir dirty output and profit; VCG remaining dirty profit rises as well. For a sufficiently small reform, the strict branch persists, giving an actual scarcity counterexample. This disproves a universal fall in protection rents. Clean replication or declining relevant demand can outweigh it, but that inequality must be established using the actual Z_j, q_j, π_j and the complete population path. Neither a low live-claim contamination rate nor a premium alone is an ecological conservation theorem.

F Paid progress, research ownership, and global institutional dynamics

This appendix supplies a connected dynamic construction. Its primitive restrictions are deliberately visible. The common public-quote economy can grow at fixed mineral expen-

diture. The permanently active nonlinear retail construction instead uses growing outside endowments and expenditure. The latter restriction prevents a growing number of retail installations from driving every installation out of the inventory regime on which its local mechanism proof rests. Neither construction treats the old CES research formula as a theorem about a changed trading institution.

F.1 Dates, ownership, scale, and the stage-market interface

Let $\delta_G, \delta_R \in (0, 1)$ be the one-period discount factors of governments and research owners, respectively. Let $g > 1$ and $L_t = g^t$. A scientific project commissioned at t operates its prototype at $t + 1$; its publicly disclosed standard becomes available for the next installation round. A successful intervention at t installs the standard disclosed at t and first produces legally at $t + 1$. Thus an already legal adopter producing at t uses $L_{t-1} = L_t/g$. The dates $-1, 0$ denote actual inherited standards, not access to an undisclosed invention. Copying a design costs no royalty; machines, site adaptation, custody and, in the nonlinear construction, authorized retail access still have paid installation contracts.

There are finitely many types of regional governments. Within type a , write $w_r > 0$ for a region's scale and μ_a^{reg} for its geographic measure. All productive capacity, household measures, intervention costs and fiscal flows of that region scale by w_r . The relevant institutional masses are

$$M_{a,t}^D = \int_{\mathcal{R}_a} w_r 1_{\{r \text{ ungoverned at } t\}} d\mu_a^{\text{reg}}, \quad M_{a,t}^C = M_a - M_{a,t}^D, \quad 0 < \sum_a M_a = \epsilon < \bar{\epsilon}.$$

These are scale-weighted geographic masses, not numbers of independent auction buyers and not masses of scientists or invention owners. A government's unilateral action is negligible; a positive fraction of governments is not. Common normalized intervention and fiscal primitives within a type make the normalized policy independent of w_r .

Fixed research-producing cores have owner masses $N_i > 0$, independent of M_a . An owner holds its own production portfolio, or an explicitly contracted royalty claim stated as part of its project. Here the nonempty construction uses direct ownership. Public adopters own no claim on a research core. If $c_0 > 0$ and $d_0 > 0$ are the initial core scales, one clean frontier and one unreformed dirty frontier give

$$\begin{aligned} B_t^C &= L_t \left(c_0 + \sum_a \frac{\omega_a^C}{g} M_{a,t}^C \right), & B_t^D &= d_0 + \sum_a \omega_a^D M_{a,t}^D, \\ \hat{B}_t^C &= B_t^C / L_t, & \hat{B}_t^D &= B_t^D / L_t. \end{aligned} \tag{186}$$

Cost coefficients and contract wedges can be absorbed in the correctly defined type capacities, but the supply and profit moments remain distinct. In a nonlinear retail cell, B in (186) denotes the mass of *complete independent installations*; its physical cost coefficient is obtained from the cell's stated normalization. In this retail construction $\hat{n}_{C,t} = \hat{B}_t^C$ and $\hat{n}_{D,t} = \hat{B}_t^D$ are the normalized clean and dirty cell masses. Replication adds the cell's

production technology, its specified finite or atomless customers, and its authorized retail access. It does not merge regional sellers or convert whole VCG buyers into Muir buyers.

Assumption F.1 (Verified stage interface for the dynamic construction). For each edition the previously proved stage market has a unique continuous allocation and actual payoff map on a compact normalized domain K . All contract, evidence and household-allocation restrictions printed in that stage theorem hold on K . Normalized prices, per-cell physical quantities, invoice receipts and net operating payoffs are bounded; the clean core has strictly positive net incremental operating value. The map includes actual fraud-organizer income, procurement, fees and the actual committed residual continuation. Historical origin and the remaining financial obligations are additional state coordinates.

For the nonlinear growth construction, normalized retail clean mass lies in $[c_0, c_0 + \epsilon \max_a \omega_a^C/g]$, normalized dirty mass in $[0, d_0 + \epsilon \max_a \omega_a^D]$, and expenditure is $E_t = L_t e_t$, with e_t in a compact positive interval. Its uncoupled retail reference price clears the genuine final-good market, with continuous cell output $Z_j(p)$:

$$p \sum_j \hat{n}_{j,t} Z_j(p) = e_t. \quad (187)$$

The left side is strictly increasing on a common price interval $[p, \bar{p}]$, with strict endpoint inequalities uniform over the state domain. The active inventory, resale and retention conditions hold throughout that interval. Zero aggregate dirty mass is permitted; the payoff of a single counterfactual dirty installation remains defined.

This is an interface to a separately solved stage problem, not an assumption that the dynamic equilibrium exists. It is directly checkable: choose the clean-only market at an interior active-premium price, take $d_0 > 0$ and $\epsilon > 0$ small, and use the strict endpoint and mechanism inequalities. Continuity gives a compact rectangle containing every state in (186), including $\hat{B}^D = 0$. Strict positivity of d_0 does not require its normalized ratio to stay away from zero. The stage counterfactuals are evaluated even for a type whose aggregate mass vanishes.

Lemma F.2 (An invariant domain for the fully connected stage). *The preceding interface is satisfied by the common public-quote growth class and by the full certification–retail coupling, with strictly positive endogenous residual bulk leakage, as follows. Write $\lambda_t = L_t^{-1}$. Safe certified capacity is $L_t b_{s,t}$, vulnerable certified capacity is a fixed $\beta_v > 0$, and the associated dirty bulk capacity $B_{D,t}$ has a fixed strictly positive core. Reform can increase b_s and reduce B_D by small scale-weighted amounts. Certified expenditure is $L_t e_c$, dirty bulk expenditure is $e_d > 0$, and retail expenditure is $L_t e_{r,t}$. The full normalized residuals are those of Theorem E.4: with $x = p_D/h$,*

$$\begin{aligned} 0 &= h^{1+\eta}[(1 + \phi)b_s + \lambda\beta_v t_v(x)] - e_c, \\ 0 &= xhY - e_d - \chi \frac{\lambda Y}{\lambda Y + \gamma Z}, \\ 0 &= pZ - e_r - \chi \frac{\lambda \gamma Z}{\lambda Y + \gamma Z}, \end{aligned} \quad (188)$$

where $Z = \sum_j \hat{n}_j Z_j(p)$, $t_v = (1 + \phi + R_v)/(1 - \pi)$, and $Y = h^\eta [B_D x^\eta - \beta_v \pi / (1 - \pi)]$. The fraud cutoff $\pi(x)$, diligence, composition losses and origin-stock constraints are the actual bulk-game quantities. For fixed small positive β_v, χ , this system has a common smooth equilibrium branch on the entire compact state rectangle and $\lambda \in [0, 1]$. Every finite date has genuinely coupled demand. Removing the retail block and setting $\chi = 0$ gives the common public-quote specialization.

Proof. At $\beta_v = \chi = 0$, the unique baseline prices are $h = [e_c / ((1 + \phi)b_s)]^{1/(1+\eta)}$, $p_D = (e_d/B_D)^{1/(1+\eta)}$, and the strict scalar retail root. Choose these prices inside the actual interior bulk fraud and active retail intervals. Positive lower bounds on b_s, B_D, Z and the strict stage derivatives make the inverse Jacobian uniformly bounded. The uniform contraction proof in Theorem E.4 then gives the same small admissible positive β_v, χ for all $\lambda \in [0, 1]$, rather than finitely many early transitions. Because vulnerable capacity stays fixed, a sufficiently small β_v has enough actual illegal-origin stock at every state. Growing safe capacity does not create an illegal origin or a new vulnerable slot. All strict stock, reporting, finance and inventory inequalities survive.

Research may own an explicitly specified complementary portfolio of safe bulk production and complete retail installations. Replication increases both parts at the same g ; its own prize sums the two genuine operating increments and their distinct paid installations and liabilities. Public adopters use the disclosed portfolio after its paid lag. Different types can use different fixed proportions. Thus $b_{s,t} = b_{s,0} + \sum_a (\omega_a^s/g) M_{a,t}^C$ and $B_{D,t} = B_{D,0} + \sum_a \omega_a^{bD} M_{a,t}^D$, while retail scales obey (186). These paths stay in the compact rectangle for every population policy. Both dirty cores and all owner masses stay strictly positive in absolute units. Retail dirty mass may tend to zero *after division by L_t* ; the bulk dirty core is not divided away. All physical throughput is $O(L_t)$ and total invoice income is $O(L_t)$. This proves the invariant-domain and funding inputs of the dynamic theorem for the complete linked class. $\square$

The fixed-expenditure throughput bound below remains a separate common benchmark. It does not assert that indefinitely growing vulnerable capacity can retain an interior fraud share when the available dirty-origin stock is fixed. The constructed infinite-path common class instead grows safe capacity and its expenditure as just specified. Within vulnerable origins, attempted admission and live contamination remain endogenous; their global weight falls as the safe flow expands. No average replaces an origin-specific stock or coalition condition.

F.2 A complete verification workload and its financing

Let Q_t be actual newly handled physical mass. Let $T_H \geq 1$ be a fixed contractual claims horizon. Evidence remains available permanently after a financial claim expires. If a unit generates at most n new authenticated records, permanent records satisfy

$$D_{t+1} = D_t + n_t Q_t, \quad 0 \leq n_t \leq n. \quad (189)$$

The existing stock is not booked again as new records. Let $a_{\text{hist}} > 0$ be the workload coefficient for unresolved historical cohorts. A conservative workload, counting physical service, unresolved vintages and retention, is

$$W_t = a_0 Q_t + a_{\text{hist}} \sum_{s=1}^{T_H} Q_{t-s} + a_D D_t, \quad a_0, a_{\text{hist}}, a_D > 0. \quad (190)$$

Missing negative-date cohorts are replaced by the actual initial ledger. This expression can overprovide service after early settlements; it never silently discards a still-open vintage. All financial and evidence queues remain separate from the physical provenance account. The finite record bound includes every permitted custody transfer and resale in the stage contract. An unbounded number of free new transactions per unit would require its own transaction workload and is excluded here.

Publicly procured verification research receives $S_V \in (0, S)$ scientists. Its paid technology is

$$Z_{t+1} = g_V Z_t, \quad g_V = 1 + \gamma_V (1 - e^{-\nu_V S_V}) > 1. \quad (191)$$

This is an exact aggregate of independent projects on an appropriately specified continuum probability space, or equivalently a deterministic research production technology. It is not a claim that a continuum eliminates common technological shocks. The public procurement contract pays the scientific wage, laboratories and testing. Deployment of each release costs $K_V > 0$ in normalized numeraire; service capital has depreciation $d_K \in (0, 1]$, paid initial installation and replacement. The productivity index Z is an efficiency of processing and storage; it is not a stock of servers that must itself be purchased unit for unit.

Proposition F.3 (Verification can keep pace with all retained obligations). *Suppose $Q_t \leq \bar{Q}G^t$, $G > 1$, $g_V \geq G$, and $Z_0 > 0$ is installed and paid for. Then W_t/Z_t is uniformly bounded, including permanent records. A fixed positive stock of service capital can serve the whole path. For every fixed $S_V \in (0, S)$ and finite $G > 1$, a finite verification productivity parameter γ_V satisfies this condition. Positive scientific resources $S - S_V$ remain for mineral research.*

Proof. The record recursion gives $D_t \leq D_0 + n\bar{Q}(G^t - 1)/(G - 1)$, plus the specified finite initial cohorts if needed. Thus, for a finite constant C_{init} ,

$$\frac{W_t}{Z_t} \leq \frac{\bar{Q}}{Z_0} \left(a_0 + a_{\text{hist}} \sum_{s=1}^{T_H} G^{-s} + \frac{a_D n}{G - 1} \right) \left(\frac{G}{g_V} \right)^t + \frac{C_{\text{init}}}{Z_0 g_V^t} =: \widehat{W}_t \leq \bar{W}.$$

Install $K \geq \bar{W}/v$ units of hardware, where $v > 0$ is service units per server-period after the Z improvement. Replacement $d_K K$ and maintenance $m_K K$ are finite every period. Acquisition $k_K K$, release deployment K_V , and research are separate positive bills. Finally $\gamma_V \geq (G - 1)/(1 - e^{-\nu_V S_V})$ is finite. This proves feasibility, not that these bills pay themselves from a zero fee. $\square$

With bounded mineral expenditure in the common power-cost class, $B_t^C = O(g^t)$ implies $Q_t = O(g^{t/(1+\eta)})$; the theorem uses $G = g^{1/(1+\eta)}$. In the nonlinear replication class prices and per-installation output remain bounded, so $Q_t = O(L_t)$ and $G = g$. If $G = 1$, permanent records grow linearly and constant Z generally fails; any $g_V > 1$ suffices instead. If storage has an irreducible positive cost per old record unaffected by Z , the permanent-record term grows without bound at fixed expenditure. That alternative technology does not satisfy the proposition.

For precise cash accounting let $c_t = 1$ in the fixed-expenditure benchmark and $c_t = L_t$ in both connected growth classes of Lemma F.2. An invoice cohort has value $I_t \leq \bar{I}_{c_t}$. Every promised loss, reward and bond is specified as a fraction of its own invoice. A conservative bond bI_t is held in its own vintage trust until its latest claim closes, where $b > 0$ covers that cohort's maximum collectible contractual liability. Different origins and vintages have different accounts. An exposed asset assigned to several claims appears only once in the joint solvency constraint. Procurement escrow is separate: an invoice dollar already due to an honest supplier is not collateral for fraud sanctions.

Proposition F.4 (Cash, capital, and resource budgets). *With finite T_H , invoice-scaled liabilities, and monotone c_t , outstanding bond principal is at most $T_H b \bar{I}_{c_t}$, apart from the known initial queue. A facility with that initial funded queue, new-cohort deposits, and explicit funding for growth and realized losses can honor all promised claims without reusing unresolved collateral. Verification, mineral research, installation, local oversight and household support admit finite uniform budget coefficients in units of c_t . A sufficiently large stated outside-good endowment and an unconditional appropriation implement them.*

Proof. At date t the cohorts still exposed belong to at most T_H different issue dates. Their principal sums to at most $b \sum_{s=0}^{T_H-1} I_{t-s} \leq T_H b \bar{I}_{c_t}$. A payout reduces only its own trust's principal; its promised maximum was funded on entry. Refunds return to the original funder at closure. Newly entering cohorts must deposit new principal, whether or not some other claim appears unlikely. An initial reserve covers actual prehistory. The facility also supplies any advance procurement escrow until the buyer's payment settles, using a different subaccount.

The resource bill consists of research scientists, laboratories, prototype manufacture, paid adoption and retail access, forensic and ordinary service, hardware installation/replacement, permanent record retention, local enforcement, and actual household transfers or consumption purchases. Each has a bounded coefficient by the workload bound, compact stage domain, finite owner and region masses, and compact action sets. A capital charge $r_B T_H b \bar{I}_{c_t}$, $r_B > 0$, explicitly prices tied bond capital; charges on procurement bridges and public liquidity reserves are added separately. Fees can fund part of these costs, with a published external appropriation funding the remainder.

For a constructive cash budget, fund each new bond deposit by an earmarked appropriation of at most $b \bar{I}_{c_t}$, return all later refunds to its funder, and never spend an anticipated refund. Initial infrastructure and the prehistory queue receive their separate date-zero appropriation. This conservative rule funds even simultaneous maximum claims; it is not

dependent on statistical diversification. A revolving facility may reduce cash advances only after applying the actual release dates.

Let $\bar{B}_{\text{real}}$ bound the resource and consumption bill and $\bar{B}_{\text{cash}}$ bound all current cash advances, including research success payments and new deposits. Specify an outside household endowment $Y_t = c_t \bar{Y}$ with $\bar{Y} > \bar{E} + \bar{B}_{\text{real}} + \bar{B}_{\text{cash}}$, plus the stated initial reserve. Buyers spend at most $c_t \bar{E}$ on mineral/final services; a lump-sum levy pays the unconditional appropriations; the remaining numeraire is strictly positive. This is a deliberately conservative feasible budget, so it does not rely on counting a transfer as a new real resource. Fees, bond principal, seizures, bribes and returned capital are reported as financial transfers. Services and forgery use real resources. Interest is a payment to capital owners and is included in private cost; the opportunity use of the tied endowment is recorded separately in the resource comparison. No sanction is simultaneously assigned to two payers. $\square$

F.3 Actual research prizes, including surviving private liabilities

Let $\mathcal{O}_{i,t+1}^s, \mathcal{O}_{i,t+1}^f$ be an owner's own operating cash receipts, net of current production and services, after success and failure. Both are evaluated at the same *actual equilibrium* next-date prices and public path: a single owner's deviation has zero mass. They include any fraud-organizer receipts actually belonging to that owner. Let K_i^z be its dated installation/access bill, B_i^z its new collateral deposit, and $C_{i,t+u}^z$ every subsequent net payment from the affected vintage, with refunds represented by negative payments. The exponent u is measured from commissioning. The same contractual claims horizon T_H closes the prototype's liabilities by $t + T_H + 1$, because its production begins at $t + 1$. The complete prize is

$$J_{i,t} = \delta_R \mathbb{E}[\Delta \mathcal{O}_{i,t+1} - \Delta K_i - \Delta B_i] - \sum_{u=2}^{T_H+1} \delta_R^u \mathbb{E} \Delta C_{i,t+u} + \delta_R^{T_H+1} \mathbb{E} \Delta V_{i,t+T_H+1}^{\text{other}}. \quad (192)$$

Here Δ means success minus failure; the final term contains any remaining difference in wealth, durable machines, ownership or operating opportunities. Committed menus and original stock stay in the affected vintage's state. A government status change never rewrites its origin.

For an owner posting its own capital, the discounted deposit and refund already value the loss of access to that capital. An imputed carry charge must not be deducted again from that same owner's prize. If a facility instead finances the deposit, the actual loan proceeds, repayments and interest replace any fictional own-capital outlay in the cash-flow ledger. The positive capital charge in Proposition F.4 is the actual funder's payment or an equivalent annuity representation, not a second charge on the same funded capital.

Lemma F.5 (When a shorter prize is valid). *Suppose temporary prototype leads expire after one operating interval, the extra prototype installations and retail rights are one-interval leases with zero subsequent salvage value (or their actual liquidation receipts are*

included in the dated cash flow), public replacement opportunities and outside financing thereafter are identical, and separate funded trusts settle every remaining liability by T_H+1 . Then $\Delta V^{\text{other}} = 0$, but the vintage-payment sum in (192) must still be retained. If a successful owner has γ_i additional units of its own baseline portfolio and operating value is linear in those units at fixed public prices, its operating increment is $\gamma_i \pi_{i,t+1}^0$. This linearity holds for the public-quote power-cost capacity and for replication of independent complete nonlinear retail cells. It need not hold for a cost change inside one nonlinear cell. Aggregate adopter surplus is not an owner's prize.

Proof. Subtract the two dated cash-flow streams. Identical later public access, funding and private opportunities cancel only after their actual remaining state differences disappear. Financial claims created by different prototype volumes do not disappear when designs are disclosed; this gives the displayed sum. In the power-cost class optimized profit equals $\beta[r]_+^{1+\eta}/(1+\eta)$, so the difference is linear in the owner's own β . For independent retail installations, maximizing the sum of n separate identical cell problems yields n times the cell value. The actual premium, retained output and selling mechanism remain inside that cell value. Changing its cost or its customer set is a different optimization and does not justify this division. Finally a non-owner has no contractual claim to the other agents' receipts. $\square$

For homogeneous core owners with aggregate own capacity B_i , each has B_i/N_i . If their common success probability is $h_i(s_i)$, aggregate next-date capacity is $B_i[1 + \gamma_i h_i(s_i)]$. Accordingly a division of *their own* aggregate operating value by $N_i[1 + \gamma_i h_i(s_i)]$ recovers an unsuccessful owner's operating value in the linear class. Legal Selva adopters never enter this numerator. With heterogeneous wedges the appropriate profit moment must be used.

The complete prize has useful primitive bounds. Write $\hat{J}_{i,t} = J_{i,t}/c_t$. On the compact market domain, suppose a clean owner's extra owned installations produce at least $o_i > 0$ and at most $\bar{o}_i$ in normalized operating income. If extra installation is at most k_i , the extra deposit at most b_i , and later extra nonrefundable payments at most ℓ_{iu} , then

$$\hat{J}_{i,t} \geq \underline{J}_i := \delta_R(o_i - k_i - b_i) - \sum_{u=2}^{T_H+1} \delta_R^u \ell_{iu}. \quad (193)$$

Refunds only improve this lower bound. An upper bound follows by retaining the maximum positive operating income and refunds and dropping costs. These are bounds on dated individual cash flows, not assumptions that an innovator obtains the entire world expenditure. They allow strictly positive institutional and installation costs. A funded success payment can cover a negative lower bound if its own appropriation is included.

Theorem F.6 (Paid verification and clean-directed mineral research). *Reserve $S_V \in (0, S)$ for the procured verification program. Mineral project owners choose effort $s_i \geq 0$, with success probability $h_i(s) = 1 - e^{-\nu_i s}$, convex nondecreasing laboratory cost $c_t R_i(s)$ with $R_i(0) = 0$, and a success payment with date- t present value $c_t G_i \geq 0$. It is actually*

paid at $t + 1$ in the amount $c_t G_i / \delta_R$, conditional on the observed success. They purchase scientists at the same normalized wage. Feasibility is exactly

$$S_V + \sum_i N_i s_{i,t} \leq S. \quad (194)$$

Choose a clean core $*$ and put $s_* = (S - S_V) / N_* > 0$. If

$$h'_*(s_*)[\underline{J}_* + G_*] - R'_*(s_*) > \max \left\{ 0, \max_{i \neq *} \left[h'_i(0)[\bar{J}_i + G_i]_+ - R'_i(0) \right] \right\}, \quad (195)$$

then every mineral research equilibrium has $s_{*,t} = s_*$ and all other $s_{i,t} = 0$. The clean standard grows endogenously at $g = 1 + \gamma_* h_*(s_*) > 1$. For any fixed stage payoff bounds and finite costs there is a finite funded G_* that makes this condition strict. It is compatible with $g_V \geq G$, positive mineral research, and $\delta_R g < 1, \delta_G g < 1$ in the connected growth classes.

Proof. An atomistic owner takes the public path and therefore its actual prize as given. For the clean core the marginal net return at every $s \in [0, s_*]$ is at least the left side of (195); it is strictly positive. For any competing project a positive net prize gives a marginal return bounded by its value at zero, while a negative prize makes every positive effort strictly worse than zero. Thus a positive unit of effort on any other project, or unused scientific labor, can be reassigned to a clean-core owner with a strict gain. Correctly weighting owners gives $N_* s_* = S - S_V$, not $s_* = S - S_V$ unless $N_* = 1$. The clean wage at s_* and actual prices supports this allocation, by strict concavity of its positive-prize objective and convex laboratory cost. Verification procures its reserved scientists at that wage, with a separate funded research contract. The technology law gives the stated g ; this is a consequence of optimal research, not an exogenous capacity injection. The aggregate grant is bounded by $c_t N_* h_*(s_*) G_* / \delta_R$ at the actual payment date. Scientific wages and research costs also have finite uniform coefficients, covered in Proposition F.4.

The coefficient of G_* on the left is strictly positive, so a finite grant works. Choose $\gamma_* > 0$ small enough that g is below the inverse discount factors, then choose finite γ_V as in Proposition F.3. The stage strict inequalities persist on a small compact rectangle. This proves a nonempty parameter class with strictly positive verification and installation costs, research in the mineral sector, both positive core masses, and $\epsilon > 0$. $\square$

Fixed outside expenditure and nonlinear replication are distinct growth normalizations. In the latter construction specify outside household endowment and buyer expenditure $Y_t = L_t \bar{Y}$, $E_t = L_t e_t$ using the g just implemented by research. This is an explicit compatible growth restriction on demand and resources. Endowments finance payments; they do not manufacture the clean standard. Holding E_t fixed while indefinitely replicating cells need not preserve an active retail regime. No result here asserts that it does.

F.4 Reclearing markets with a positive mass of governments

Each ungoverned government chooses between waiting and one feasible positive intervention intensity. Its normalized all-in cost $k_a > 0$ includes force, preparation/installation, advance funding cost, and any temporary household-support obligation; success probability $P_a \in (0, 1)$ is strictly positive. A finite menu of intensities changes the argument only by replacing a binary probability by a probability simplex. An action-independent outside appropriation or initially endowed facility supplies liquidity; its common flow cancels from the political comparison. Spending k_a still sacrifices numeraire value. An action-contingent reimbursement would alter that comparison and is not silently omitted.

Let $R_{a,t}$ be legal tax receipts less actual local services and support, and $b_{a,t}$ the actual protection payment drawn from the illegal actor's distributable net receipts, including its own net fraud income. Write $\hat{R}_{a,t} = R_{a,t}/(w_r c_t)$ and $\hat{b}_{a,t} = b_{a,t}/(w_r c_t)$. Profit tax $\tau_a \in (0, 1)$ is levied on the optimized legal operating surplus; a proportional tax preserves that static optimization. Paid adoption is already included in the promised lag and the installation budget. In the nonlinear growth class (and its growing common specialization) $c_t = L_t$ and the normalized government continuation factor is $\hat{\delta}_G = \delta_G g < 1$; at fixed expenditure $c_t = 1$ and $\hat{\delta}_G = \delta_G < 1$. Normalized values obey

$$\begin{aligned} U_{a,t} &= \hat{R}_{a,t} + \hat{\delta}_G U_{a,t+1}, \\ V_{a,t} &= \max\{\hat{b}_{a,t} + \hat{\delta}_G V_{a,t+1}, -k_a + \hat{\delta}_G [P_a U_{a,t+1} + (1 - P_a) V_{a,t+1}]\}. \end{aligned} \quad (196)$$

The two institutional states are sufficient because historical claims remain obligations of their original signers and their funded vintage trusts. The incoming government and legal operator do not assume an outgoing organizer's debts; the trusts remain enforceable after reform. Current legal fees include the full prospective cost of their own new cohorts. A state's past tenure therefore creates no omitted public debt or treasury state. Waiting and failed intervention leave the same current production and funded vintage obligations. If a government guarantees old claims, or if finance depends on its past success, those private or public debt coordinates must instead enter its Bellman equation.

The current production allocation is unchanged until the next date's success: intervention forfeits the current protection payment but does not invent a separate current illegal-output equation. If production disruption is modeled, its actual cost and supply effect must be added.

If $u_{a,t} \in [0, 1]$ is the fraction choosing intervention, the exact population update is

$$\begin{aligned} \sigma_{a,t} &= u_{a,t} P_a, \\ M_{a,t+1}^D &= (1 - \sigma_{a,t}) M_{a,t}^D, & M_{a,t+1}^C &= M_{a,t}^C + \sigma_{a,t} M_{a,t}^D. \end{aligned} \quad (197)$$

An exact law of large numbers for independent regional outcomes is part of this deterministic atomless construction. With a common intervention shock the distribution is a state and the deterministic law cannot be used. After every transition (186) changes, the market clears again, all own research prizes are repriced, and government continuations

are recomputed. Genuine indifference permits mixing; away from it the chosen action is pure.

Theorem F.7 (Connected infinite-horizon equilibrium). *Suppose the verified stage interface, the funded workload and budget conditions, and the strict research-direction condition hold. Suppose normalized local flows and vintage obligations are continuous and bounded on the stated compact domain. In either growing-income class require $\max(\delta_G, \delta_R)g < 1$. There exists an infinite-horizon equilibrium of the stage market, science allocation, funded verification, government policies, and positive-mass population transitions. Government choices include waiting and intervention, with mixing only at equality of their continuation values. Initial historical obligations are honored. The theorem makes no global uniqueness claim.*

Proof. The strict research inequality fixes the mineral effort vector at every candidate market path; it does not fix its prices or prizes. Verification research fixes its paid productivity transition. For a sequence of government fractions $u = (u_{a,t})$, the population recursion and adoption lag give a unique normalized capacity path. The stage theorem determines prices and the actual current allocation at each date. The finite claims queue, normalized permanent-record stock, and reserve ledgers update continuously from the specified pre-history. For instance $D_{t+1}/Z_{t+1} = (D_t/Z_t + n_t Q_t/Z_t)/g_V$, so the permanent record state is bounded, although unnormalized records are permanent and unbounded. The remaining finite vintage coordinates are bounded after division by their stated invoice/scale units. Thus every generated path lies in a fixed compact product of finite-dimensional state sets.

First truncate at T , set terminal normalized government values to zero, and keep actual finite contractual settlements through their closing dates when computing research prizes. Research can equivalently be continued at its strictly dominant allocation past T ; only government terminal values are truncated. On the cube $[0, 1]^{A(T+1)}$, a candidate u generates a continuous finite public path. Backward induction in (196) gives continuous values and continuous intervention advantages. The best-response correspondence at (a, t) equals $\{0\}$ for a negative advantage, $\{1\}$ for a positive advantage, and $[0, 1]$ at zero. It has nonempty compact convex values and closed graph. Their finite product has a fixed point by the finite-dimensional Kakutani theorem [Kakutani, 1941]. Unilateral deviations keep the aggregate path fixed, as required for atomless governments.

Choose such a fixed point for each T and extend its fractions arbitrarily beyond T . A diagonal subsequence converges coordinatewise at every finite date. The forward recursions and continuity of the stage map imply convergence of all finite-date states, prices, allocations and liabilities. If the normalized absolute one-period political payoff is bounded by $\bar{r}$, changing every continuation after date T changes a date- t value by at most $2\bar{r}\hat{\delta}_G^{T-t+1}/(1 - \hat{\delta}_G)$, uniformly in the population path and in the chosen policies. For research, the finite liability window and any explicitly retained continuation have the corresponding uniform bound with $\delta_R g$ (or δ_R in the conditional fixed-expenditure benchmark). These tails tend uniformly to zero. Therefore limiting finite-date advantages equal the advantages in the true infinite continuation. If a limiting intervention fraction

is positive where its advantage is strictly negative, continuity and the tail bound give the same strict negative sign in all sufficiently long truncations, contradicting their fixed-point condition. The waiting case is identical. Equality supports the limiting mixture.

Passing to the limit in the population recursion and all ledgers preserves their identities, nonnegative balances, and initial conditions. The mineral research allocation remains privately optimal because its uniform strict inequality was evaluated against bounds covering these actual prices and liabilities. Its actual success–failure values are given by (192), not by the truncated guess. Every budget remains funded by its uniform appropriation. This proves a connected equilibrium. Discounting was used to bound tails and to solve a Bellman problem on a given path. It was never used to declare the full market/research/population map a contraction. $\square$

F.5 Scarcity can strengthen protection: counterexample and re-pair

Proposition F.8 (Scarcity and a sufficient dominance condition). *At a scalar final-good market of the form (187), reducing dirty capacity at fixed clean capacity and demand strictly raises price whenever the removed capacity has positive final output. If an illegal cell’s optimized payoff is strictly increasing in price, the protection rent of each remaining such cell rises. A fall in total illegal production therefore does not imply a fall in every remaining region’s rent.*

Conversely, for two consecutive normalized states, price does not rise if at the old price p_t

$$p_t \sum_j (\hat{n}_{j,t+1} - \hat{n}_{j,t}) Z_j(p_t) \geq e_{t+1} - e_t. \quad (198)$$

Strict inequality gives a strict price fall. In the explicit nonlinear growth class take $e_t = e_0 + a\rho^t$, $g^{-1} \leq \rho < 1$. If

$$a(1 - \rho) > (d_0 + \epsilon \max_a \omega_a^D) \max_{p \in [\underline{p}, \bar{p}]} p Z_D(p), \quad (199)$$

then prices fall at every transition, for every admissible reform path. Clean technology grows through its optimally funded research; this additional declining normalized expenditure policy makes its supply and demand effects dominate the loss of dirty supply.

Proof. Let $F(p) = p \sum_j \hat{n}_j Z_j(p) - e$. It is strictly increasing and vanishes at the old equilibrium price. Removing positive dirty mass makes its value there negative, so the new root is higher. The envelope theorem for a cell with objective $p\Psi - C$ gives $\pi'(p) = \Psi > 0$ at the active optimum; the genuine premium and retained material remain in Ψ . For the converse, (198) says that the new excess supply evaluated at the old root is nonnegative. Strict monotonicity of F gives the claimed root order without differentiating a policy path.

In (186) normalized clean mass is nondecreasing and normalized dirty mass satisfies $0 \leq \hat{n}_{D,t} - \hat{n}_{D,t+1} \leq (d_0 + \epsilon \max \omega_a^D) g^{-t}$, even if every ungoverned region intervenes. The normalized expenditure fall is $a(1 - \rho)\rho^t$, at least that coefficient times g^{-t} . Condition (199) therefore strictly covers the largest possible loss of normalized dirty output at every old equilibrium price. The increase of clean capacity from successful adoption only strengthens the comparison. Choose d_0, ϵ small and strictly positive; then the required a is small enough to preserve the strict stage price interval. Thus this is a nonempty primitive condition on supplies and a funded expenditure path, not a restatement that equilibrium rents decline. $\square$

Corollary F.9 (Dominance in the fully coupled infinite path). *Use Lemma F.2. Let every reformable type install a strictly positive retail-clean component, so changes in b_s and B_D satisfy $|\Delta b_s| + |\Delta B_D| \leq A\Delta\hat{n}_C$ for a finite primitive ratio A . Keep e_c, e_d constant and let $e_{r,t} = e_0 + a\rho^t$, $g^{-1} \leq \rho < 1$. There are finite uniform positive derivative bounds m_C, m_E, M_D, C obtained from the strict baseline market and Theorem E.4. If*

$$C\chi A < m_C, \quad m_E a(1 - \rho) > M_D(d_0 + \epsilon \max \omega_a^D) + C\chi(1 - g^{-1}), \quad (200)$$

the actual retail price in the full three-price equilibrium strictly decreases along every admissible government path. The class is nonempty with all masses, β_v, χ , and paid institutional costs strictly positive. The conclusion concerns the actual retail dirty regions; it does not assert that the auxiliary bulk dirty price also falls under this same one-component demand policy.

Proof. On the uncoupled compact domain the price derivatives satisfy $p_{n_C}^0 < 0$, $p_{e_r}^0 > 0$, and $p_{n_D}^0 < 0$, uniformly away from zero; this follows by differentiating the strictly increasing scalar clearing equation. Its derivatives with respect to b_s, B_D, λ are zero. Choose half of the positive derivative margins for m_C, m_E and a finite upper bound for M_D . The uniform C^1 estimate in Theorem E.4, with χ small, preserves $p_{n_C} \leq -m_C$, $p_{e_r} \geq m_E$, $|p_{n_D}| \leq M_D$ and $|p_{b_s}| + |p_{B_D}| + |p_\lambda| \leq C\chi$ after enlarging C . The derivative estimates hold on a convex rectangle, so integrate them along the straight segment between two successive states. It follows that

$$\Delta p \leq -(m_C - C\chi A)\Delta\hat{n}_C + M_D|\Delta\hat{n}_D| - m_E|\Delta e_r| + C\chi|\Delta\lambda|.$$

Use $|\Delta\hat{n}_D| \leq (d_0 + \epsilon \max \omega_a^D)g^{-t}$, $|\Delta\lambda| = (1 - g^{-1})g^{-t}$ and $|\Delta e_r| = a(1 - \rho)\rho^t$. Conditions (200) give a strict negative sign at every date. There is no constant perturbation error competing with a vanishing drift. Choose small positive retail dirty core and reformable mass, then a small positive demand drift a within the stage interval, then positive χ small enough for both strict inequalities. The fixed bulk core can remain positive. This produces the required open nonempty parameter class. $\square$

A transparent counterexample uses $\eta = 1$, a dirty supply coefficient B , and fixed dirty expenditure $E_D = 1$. Clearing yields $p_D = B^{-1/2}$. A remaining unit-scale region produces p_D and earns $p_D^2/2 = 1/(2B)$. Reducing B from 1 to 0.81 raises price from 1 to 10/9,

that region's output by 11.1%, and its profit from 0.5 to 50/81. Total dirty output falls from 1 to 0.9. This is an exact regional optimum and market-clearing counterexample. In the active nonlinear market the preceding root and envelope proof gives the same rent counterexample with its actual mechanism, without adding a fictitious monopoly term to a public-quote formula.

For physical extraction the relevant cell function is $q_D(p)$, not the transformed final-good quantity $Z_D(p)$. In the proved nonlinear Muir branch $q'_D(p) > 0$; in the stated finite VCG inventory branch q_D is constant, so the weak sign holds. If this physical monotonicity is retained on the compact domain, (198), or the full coupled condition (200), implies that each unchanged dirty retail cell's extraction weakly falls and total reformable retail dirty extraction weakly falls. A region owning an additional bulk dirty component need not experience the same sign for its combined extraction or protection rent. An inequality for final output alone would not prove that claim. The fixed dirty core may continue to extract forever.

F.6 Intervention without assuming a monotone political scissor

Theorem F.10 (Eventually strict intervention on the recleared path). *In any connected equilibrium suppose each type has a funded legal-flow floor $\hat{R}_{a,t} \geq \underline{R}_a > 0$ and normalized protection flows satisfy $0 \leq \hat{b}_{a,t} \rightarrow 0$ uniformly over feasible population paths. If*

$$k_a(1 - \hat{\delta}_G) < \hat{\delta}_G P_a \underline{R}_a, \quad (201)$$

there is a finite date after which every still-ungoverned government of that type strictly chooses intervention. Its remaining scale mass then decays by $1 - P_a$ per date. An individual government's success time is almost surely finite. This does not require R_t or b_t to be monotone on the world path.

Proof. For a fixed actual public path the normalized Bellman equations have unique bounded values. Subtract them and write $D_{a,t} = U_{a,t} - V_{a,t}$:

$$D_{a,t} = \min\{\hat{R}_{a,t} - \hat{b}_{a,t} + \hat{\delta}_G D_{a,t+1}, \hat{R}_{a,t} + k_a + \hat{\delta}_G(1 - P_a)D_{a,t+1}\}.$$

If $\hat{b}_{a,t} \leq b_*$ for all $t \geq T$, comparison of this increasing contraction with its constant-flow version gives

$$D_{a,t} \geq D_* := \min\left\{\frac{\underline{R}_a - b_*}{1 - \hat{\delta}_G}, \frac{\underline{R}_a + k_a}{1 - \hat{\delta}_G(1 - P_a)}\right\}, \quad t \geq T.$$

To check this comparison without a terminal assumption, iterate both operators from the same bounded sequence; monotonicity preserves the order and contraction removes the initial difference. The intervention advantage relative to waiting is $-k_a - \hat{b}_{a,t} + \hat{\delta}_G P_a D_{a,t+1}$, at least $-k_a - b_* + \hat{\delta}_G P_a D_*$. At $b_* = 0$ this lower bound is strictly positive exactly when (201) holds (check both fractions). By continuity it remains positive for some

$b_* > 0$. Uniform disappearance of normalized bribes gives the claimed finite T . Population recursion then gives geometric decay. For a particular region, the probability of n further failures is $(1 - P_a)^n \rightarrow 0$. $\square$

In both connected growing-expenditure constructions an unreformed region retains fixed dirty capacity, bounded prices and bounded actual fraud opportunities, so $b_{a,t}/w_r \leq \bar{b}_a$ even if scarcity raises its rent initially. Dividing by L_t makes its normalized protection flow tend uniformly to zero. A legal region installs L_t/g units per scale, so a strictly positive per-cell tax base can fund a strictly positive normalized $\underline{R}_a$. Thus the theorem has direct primitive content. Choose $0 < \tau_a < 1$, positive but sufficiently small local service and support costs, and a strictly positive k_a satisfying (201). All inequalities are open. The optional demand-drift condition strengthens this result to falling *absolute* dirty rents; growing legal opportunity already suffices for the normalized crossing.

For a fixed-expenditure common market, bounded clean operating receipts give a constant legal-flow floor after proportional adoption and its lag. The same proof applies if the actual dirty protection flow tends to zero, as under a separately proved sufficiently strong clean substitution and demand transition. It does not apply to a fixed dirty expenditure share that leaves protection rents positive. More generally the proof gives a uniform intervention test directly with an upper bound $b_* > 0$ through $-k - b_* + \hat{\delta}_G PD_* > 0$. If this bound fails, waiting can be optimal forever. For example $R_t \equiv R$, $b_t \equiv b \geq R$ implies perpetual waiting when $k > 0$: the value of immediate legal access cannot compensate for forfeiting at least the same stream and paying the intervention cost. If $P = 0$, intervention also has no benefit. These are retained high-cost/non-intervention cases, not failures of the existence theorem.

Waiting can be strictly optimal before intervention becomes attractive. For a primitive two-phase example take $\hat{\delta}_G = 0.8$, $P = 0.5$, $k = 0.1$, $R_t \equiv 0.2$, and $b_t = 0$ for $t \geq 1$. Then the stationary later solution has $U = 1$, $V = 0.5$, and intervention advantage $0.1 > 0$. At date zero set $b_0 = 0.2$. Using the actual next-date values gives waiting value 0.6 and intervention value 0.5. Hence it strictly waits once, then strictly intervenes. A small continuous change of the initial demand/protection state preserves both strict inequalities; mixing occurs only on their equality boundary.

F.7 Returning the theorem to households and a physical forest

Each output in this appendix is reconstructed with the regional implementation correspondence Γ_r . Let $y_{rh}(q)$ be a household's actual feasible net livelihood on the selected legal branch, after effort and its baseline alternatives are valued in the stated numeraire convention. For a promised floor $\underline{y}_{rh}$ set

$$T_{rh,t} = [\underline{y}_{rh} - y_{rh}(q_{r,t})]_+. \quad (202)$$

An outside-good transfer does not change its mining first-order condition when it is a committed eligibility-based floor payment independent of the household's subsequent marginal

output choice. Operationally the payment is fixed from the committed implemented allocation and is not withdrawn dollar for dollar after an unobserved earnings deviation; such a withdrawal rule would require solving a new household incentive problem. Assume the proved constructive household branch and the announced support contract give a continuous aggregate transfer with $\int T_{rh,t} dh \leq w_r c_t f_{L,a}$. This bound follows, for example, from a bounded per-scale household measure and bounded livelihood shortfalls. It is explicitly charged to local taxes or the unconditional outside appropriation. It is not part of cartel profit merely because household rent appeared in procurement cost.

If the legal per-scale operating surplus satisfies $\pi_{a,t}^C/(w_r c_t) \geq \underline{\pi}_a > 0$, local services cost at most $w_r c_t f_{S,a}$, and support at most $w_r c_t f_{L,a}$, the primitive funding condition is

$$0 < f_{S,a} + f_{L,a} < \tau_a \underline{\pi}_a, \quad \underline{R}_a = \tau_a \underline{\pi}_a - f_{S,a} - f_{L,a} > 0. \quad (203)$$

A separate advance reserve covers one local settlement period. Initial adoption, retail rights where needed, continuing replacement, and household support during failed interventions are included in the specified public or operator accounts. The government is not lending money that only exists after an unfinanced success.

Proposition F.11 (What governance does and does not conserve). *Let the critical forest contain the reformable Selva population, while the fixed research cores operate outside that critical habitat. Require the legal implementation set to enforce zero extraction disturbance in the protected forest (for example, paid legal installations use a separately specified noncritical site). Let an ungoverned region's disturbance be at most $w_r \bar{d}_a$ on the compact physical allocation domain. Then eventual intervention implies the population's expected or exact continuum aggregate critical disturbance tends to zero. Every strictly positive tolerance is reached after a finite date. Zero disturbance at a common finite date is not implied for an atomless population with $P_a < 1$.*

If no irreversible loss is permitted, a cumulative-disturbance budget $\mathcal{B}$ is sufficient only if it covers the transition as well: for a common intervention-start bound T ,

$$\sum_a \bar{d}_a M_{a,0}^D \left(T + \frac{1}{P_a} \right) \leq \mathcal{B}. \quad (204)$$

Successful legal production must itself remain in the specified safe implementation set. A certificate is not such a physical restriction.

Proof. Before T disturbance is bounded by $\sum_a \bar{d}_a M_{a,0}^D$ each period. After T the remaining masses are at most $M_{a,0}^D (1 - P_a)^{t-T}$. Their disturbance bounds tend to zero and their geometric sums equal $M_{a,0}^D / P_a$. This proves the tolerance and cumulative bounds. At every finite date, however, $M_{a,0}^D (1 - P_a)^{t-T} > 0$ when the initial mass is positive and $P_a < 1$. Almost-surely finite success for each one of a continuum of regions does not provide a common finite success date. For finitely many separately designated critical regions, the maximum of their almost-surely finite success times is finite almost surely, but a finite deadline still has positive failure probability. Permanent historical evidence

and illegal origin remain in the ledger after all these governance changes; revocation changes status, not matter. $\square$

The conclusion is therefore about a financed transition into a specified safe technology and feasible household livelihoods. It is not a theorem that lower live certificate contamination by itself lowers physical extraction, that every illegal supplier disappears, or that institutional success can reverse damage already incurred before control was established.

References

- Ryan Abman. Rule of law and avoided deforestation from protected areas. *Ecological Economics*, 146:282–289, 2018.
- Daron Acemoglu. Directed technical change. *Review of Economic Studies*, 69(4): 781–809, 2002.
- Daron Acemoglu. Why not a political coase theorem? social conflict, commitment, and politics. *Journal of Comparative Economics*, 31(4):620–652, 2003.
- Daron Acemoglu. Politics and economics in weak and strong states. *Journal of Monetary Economics*, 52(7):1199–1226, 2005. doi: 10.1016/j.jmoneco.2005.05.001. URL <https://doi.org/10.1016/j.jmoneco.2005.05.001>.
- Daron Acemoglu. Modeling inefficient institutions. In Richard Blundell, Whitney K. Newey, and Torsten Persson, editors, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, volume 1, chapter 9, pages 341–380. Cambridge University Press, 2006. doi: 10.1017/CBO9781139052269.011.
- Daron Acemoglu and James A. Robinson. Political losers as a barrier to economic development. *American Economic Review*, 90(2):126–130, 2000.
- Daron Acemoglu and James A. Robinson. Economics versus politics: Pitfalls of policy advice. *Journal of Economic Perspectives*, 27(2):173–192, 2013. doi: 10.1257/jep.27.2.173. URL <https://www.aeaweb.org/articles?id=10.1257/jep.27.2.173>.
- Daron Acemoglu, Pol Antràs, and Elhanan Helpman. Contracts and technology adoption. *American Economic Review*, 97(3):916–943, 2007.
- Daron Acemoglu, Davide Ticchi, and Andrea Vindigni. Emergence and persistence of inefficient states. *Journal of the European Economic Association*, 9(2):177–208, 2011.
- Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American Economic Review*, 102(1): 131–166, 2012. doi: 10.1257/aer.102.1.131.
- Daron Acemoglu, James A. Robinson, and Rafael J. Santos. The monopoly of violence: Evidence from colombia. *Journal of the European Economic Association*, 11(S1):5–44, 2013.
- Daron Acemoglu, Camilo García-Jimeno, and James A. Robinson. State capacity and economic development: A network approach. *American Economic Review*, 105(8): 2364–2409, 2015. doi: 10.1257/aer.20140044. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20140044>.
- Daron Acemoglu, Ufuk Akcigit, Douglas Hanley, and William Kerr. Transition to clean technology. *Journal of Political Economy*, 124(1):52–104, 2016.

- Rafael Araujo. When clouds go dry: An integrated model of deforestation, rainfall, and agriculture. *Working Paper*, 2024.
- Rafael Araujo, Juliano Assunção, Marina Hirota, and José A. Scheinkman. Estimating the spatial amplification of damage caused by degradation in the amazon. *Proceedings of the National Academy of Sciences*, 120(46):e2312451120, 2023.
- Rafael Araujo, Francisco Costa, and Marcelo Sant’Anna. Efficient conservation of the Brazilian Amazon: Estimates from a dynamic model. Working paper, version of 10 December 2024, 2024.
- Juliano Assunção, Clarissa Gandour, and Romero Rocha. DETER-ing deforestation in the amazon: Environmental monitoring and law enforcement. *American Economic Journal: Applied Economics*, 15(2):125–156, 2023.
- Juliano Assunção, Lars Peter Hansen, Todd Munson, and José A. Scheinkman. Carbon prices and forest preservation over space and time in the Brazilian Amazon. Technical Report 2023-86, Becker Friedman Institute, 2025. Revised 30 June 2025.
- Lawrence M. Ausubel and Peter Cramton. The optimality of being efficient. Working paper, University of Maryland, June 1999. URL <https://ausubel.com/auction-papers/99wp-optimality-of-being-efficient.pdf>. Version dated 18 June 1999.
- Clare Balboni, Aaron Berman, Robin Burgess, and Benjamin A. Olken. The economics of tropical deforestation. *Annual Review of Economics*, 15:723–754, 2023. doi: 10.1146/annurev-economics-090622-024705.
- Mark Banchereau. Rebel group in Congo generates \$300,000 monthly in seized mining area, UN says. Associated Press, 2024. URL <https://abcnews.go.com/International/wireStory/rebel-group-congo-generates-300000-monthly-seized-mining-114377360>.
- Kathryn Baragwanath and Ella Bayi. Collective property rights reduce deforestation in the brazilian amazon. *Proceedings of the National Academy of Sciences*, 117(34): 20495–20502, 2020.
- Björn Bartling, Roberto A. Weber, and Lan Yao. Do markets erode social responsibility? *Quarterly Journal of Economics*, 130(1):219–266, 2015. doi: 10.1093/qje/qju031.
- Panle Jia Barwick, Shanjun Li, Ligu Lin, and Eric Yongchen Zou. From fog to smog: The value of pollution information. *American Economic Review*, 114(5):1338–1381, 2024. doi: 10.1257/aer.20200956.
- Kaushik Basu and Pham Hoang Van. The economics of child labor. *American Economic Review*, 88(3):412–427, 1998. URL https://qed.econ.queensu.ca/pub/faculty/sumon/basu_childlabor.pdf.

- Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019. doi: 10.1257/jel.20181489.
- Nicolas Berman, Mathieu Couttenier, Dominic Rohner, and Mathias Thoenig. This mine is mine! how minerals fuel conflicts in Africa. *American Economic Review*, 107(6):1564–1610, 2017. doi: 10.1257/aer.20150774. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20150774>.
- 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>.
- Bloomberg News. Mining murders show Peru’s struggles with illegal gold rush. Bloomberg News, republished by MINING.COM, May 2025. URL <https://www.mining.com/web/mining-murders-show-perus-struggles-with-illegal-gold-rush/>. SNMPE describes REINFO as a shield used by illegal miners attacking formal operations.
- Leonardo Bonilla-Mejía and Iván Higuera-Mendieta. Protected areas under weak institutions: Evidence from colombia. *World Development*, 122:585–596, 2019.
- Chris A Boulton, Timothy M Lenton, and Niklas Boers. Pronounced loss of amazon rainforest resilience since the early 2000s. *Nature Climate Change*, 12(3):271–278, 2022.
- A. Lans Bovenberg and Lawrence H. Goulder. Environmental taxation and regulation. In Alan J. Auerbach and Martin Feldstein, editors, *Handbook of Public Economics*, volume 3, chapter 23, pages 1471–1545. Elsevier, 2002.
- Avishay Braverman and Joseph E. Stiglitz. Sharecropping and the interlinking of agrarian markets. *American Economic Review*, 72(4):695–715, 1982. URL <https://documents1.worldbank.org/curated/en/975871468344645509/pdf/REP229000Share00of0agrarian0markets.pdf>.
- Robin Burgess, Matthew Hansen, Benjamin A Olken, Peter Potapov, and Stefanie Sieber. The political economy of deforestation in the tropics. *The quarterly journal of economics*, 127(4):1707–1754, 2012.
- Gilles Chemla. Downstream competition, foreclosure, and vertical integration. *Journal of Economics & Management Strategy*, 12(2):261–289, 2003.
- Nicola Comincioli, Christopher R. Knittel, Elsa A. Olivetti, Ilenia G. Romani, and Sergio Vergalli. Critical minerals for a just energy transition. *Annual Review of Resource Economics*, 17:241–262, 2025. doi: 10.1146/annurev-resource-011224-050431.
- Brian R. Copeland and M. Scott Taylor. Trade, growth, and the environment. *Journal of Economic Literature*, 42(1):7–71, 2004.

- Partha Dasgupta and Debraj Ray. Inequality as a determinant of malnutrition and unemployment: Theory. *Economic Journal*, 96(384):1011–1034, 1986. doi: 10.2307/2233171. URL <https://debrajray.com/wp-content/uploads/2017/12/DasguptaRay86.pdf>.
- Oeindrila Dube and Juan F. Vargas. Commodity price shocks and civil conflict: Evidence from Colombia. *Review of Economic Studies*, 80(4):1384–1421, 2013. doi: 10.1093/restud/rdt009. URL <https://doi.org/10.1093/restud/rdt009>.
- Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *Quarterly Journal of Economics*, 128(4):1499–1545, 2013. doi: 10.1093/qje/qjt024. URL <https://doi.org/10.1093/qje/qjt024>.
- P. Dunia Kabunga and Sara Geenen. Work regimes and technology-driven transformations in DR Congo’s artisanal and small-scale gold mines. *The Extractive Industries and Society*, page 101142, 2022. doi: 10.1016/j.exis.2022.101142. URL <https://medialibrary.uantwerpen.be/files/1253738/c5e9627d-0c99-4e4b-a19f-cbc89a4eeb5b.pdf>.
- European Parliament and Council of the European Union. Regulation (EU) 2017/821 laying down supply chain due diligence obligations for union importers of tin, tantalum and tungsten, their ores, and gold originating from conflict-affected and high-risk areas. Official Journal of the European Union, L 130, 19 May 2017, May 2017. Adopted 17 May 2017; due diligence obligations apply from 1 January 2021.
- Benjamin Faber, Benjamin Krause, and Raúl Sánchez de la Sierra. Artisanal mining, livelihoods, and child labor in the cobalt supply chain of the democratic republic of congo. Policy report, Center for Effective Global Action, University of California, Berkeley, May 2017. URL <https://escholarship.org/uc/item/17m9g4wm>.
- Farid Farrokhi and Ahmad Lashkaripour. Can trade policy mitigate climate change? *Econometrica*, 93(5):1561–1599, 2025.
- Farid Farrokhi, Elliot Kang, Heitor S. Pellegrina, and Sebastian Sotelo. Deforestation: A global and dynamic perspective. Working paper, December 2025 version, 2025.
- Bernardo M. Flores, Encarni Montoya, Boris Sakschewski, Nathalia Nascimento, Arie Staal, Richard A. Betts, Carolina Levis, David M. Lapola, Adriane Esquivel-Muelbert, Catarina Jakovac, et al. Critical transitions in the amazon forest system. *Nature*, 626(7999):555–564, 2024.
- GS1. *EPCIS and CBV Implementation Guideline*. GS1, 2.0 edition, 2023. URL <https://ref.gs1.org/guidelines/epcis-cbv/2.0.0/>. Accessed 8 September 2026.
- Hanna Halaburda. Blockchain revolution without the blockchain? *Communications of the ACM*, 61(7):27–29, 2018.

- Robert Heinkel, Alan Kraus, and Josef Zechner. The effect of green investment on corporate behavior. *Journal of Financial and Quantitative Analysis*, 36(4):431–449, 2001.
- Soeren J. Henn, Christian Mastaki Mugaruka, Miguel Ortiz, Raúl Sánchez de la Sierra, and David Qihang Wu. Monopoly of taxation without a monopoly of violence: The weak state’s trade-offs from taxation. *Review of Economic Studies*, 92(2):1126–1156, 2025. doi: 10.1093/restud/rdae036.
- Jeffrey I. Herbst. *States and Power in Africa: Comparative Lessons in Authority and Control*. Princeton University Press, Princeton, NJ, 2000.
- International Energy Agency. Global critical minerals outlook 2025. Technical report, International Energy Agency, 2025. URL <https://www.iea.org/reports/global-critical-minerals-outlook-2025/executive-summary>. Executive summary; accessed 8 September 2026.
- Seema Jayachandran, Joost De Laat, Eric F Lambin, Charlotte Y Stanton, Robin Audy, and Nancy E Thomas. Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science*, 357(6348):267–273, 2017.
- Shizuo Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Mathematical Journal*, 8(3):457–459, 1941. doi: 10.1215/S0012-7094-41-00838-4. URL <https://www.math.ucdavis.edu/~deloera/TEACHING/READINGSEMINAR/PAPERS/kakutanioriginal.pdf>.
- Rachel E. Kranton and Deborah F. Minehart. A theory of buyer-seller networks. *American Economic Review*, 91(3):485–508, 2001. doi: 10.1257/aer.91.3.485.
- Vijay Krishna and Motty Perry. Efficient mechanism design. Technical Report 133, Center for Rationality and Interactive Decision Theory, Hebrew University of Jerusalem, 1998. URL <https://ratio.huji.ac.il/files/dp133.pdf>. Version dated 14 April 1998.
- David M Lapola, Patricia Pinho, Jos Barlow, Luiz EOC Aragão, Erika Berenguer, Rachel Carmenta, Hannah M Liddy, Hugo Seixas, Camila VJ Silva, Celso HL Silva-Junior, et al. The drivers and impacts of amazon forest degradation. *Science*, 379(6630):eabp8622, 2023.
- Arik Levinson and M. Scott Taylor. Unmasking the pollution haven effect. *International Economic Review*, 49(1):223–254, 2008.
- 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>.

- Simon Loertscher and Ellen V. Muir. Market power, randomization and regulation. *International Journal of Industrial Organization*, 96:103081, 2024. doi: 10.1016/j.ijindorg.2024.103081. URL https://simonloertscher.net/wp-content/uploads/2024/05/MPRR_Loertscher.Muir_.2024_05_22.pdf. Author version dated 22 May 2024.
- London Bullion Market Association. Responsible gold guidance. LBMA, London, 2012. Core of the LBMA Responsible Sourcing Programme; compliance required for Good Delivery List refiners.
- London Bullion Market Association. About good delivery, 2026. URL <https://www.lbma.org.uk/good-delivery/about-good-delivery>. Accessed 8 September 2026.
- London Metal Exchange. Trading, 2026. URL <https://www.lme.com/Trading>. Accessed 8 September 2026.
- Mihai Manea. Intermediation and resale in networks. *Journal of Political Economy*, 126(3):1250–1301, 2018. doi: 10.1086/697205.
- Mihai Manea. Bottleneck links, essential intermediaries, and competing paths of diffusion in networks. *Theoretical Economics*, 16(3):1017–1053, 2021. doi: 10.3982/TE4385.
- Robert McDonald and Daniel Siegel. The value of waiting to invest. *Quarterly Journal of Economics*, 101(4):707–728, 1986. doi: 10.2307/1884175.
- Paul R. Milgrom, Douglass C. North, and Barry R. Weingast. The role of institutions in the revival of trade: the law merchant, private judges, and the champagne fairs. *Economics & Politics*, 2(1):1–23, 1990.
- Luc Moreau and Paolo Missier. PROV-DM: The PROV data model. W3c recommendation, World Wide Web Consortium, April 2013. URL <https://www.w3.org/TR/prov-dm/>. Recommendation of 30 April 2013; accessed 8 September 2026.
- Paramita Mukherjee and Vivekananda Mukherjee. Assessing the scope of gold refineries in india: Evidence from survey data. *Resources Policy*, 69:101810, 2020. doi: 10.1016/j.resourpol.2020.101810.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981. doi: 10.1287/moor.6.1.58.
- OECD. *OECD Due Diligence Guidance for Responsible Supply Chains of Minerals from Conflict-Affected and High-Risk Areas: Third Edition*. OECD Publishing, Paris, 2016. doi: 10.1787/9789264252479-en. URL <https://doi.org/10.1787/9789264252479-en>.

- OECD. Costs and value of due diligence in mineral supply chains. Oecd position paper, OECD Publishing, Paris, 2021. URL <https://doi.org/10.1787/4d432567-en>.
- OECD and UN. Conflict transformation and the role of responsible artisanal and small-scale mining: Supporting peace through supply chain due diligence, 2023. URL <https://doi.org/10.1787/cdbd61d1-en>.
- R. Pacsi, S. Novoa, O. Yupanqui, M. Quispe, S. La Torre, H. Balbuena, B. Huamán, G. Valdivia, C. Castañeda, M. Soria, M. Finer, and A. Santana. Current situation of gold mining in the Peruvian Amazon. MAAP Report 233, Monitoring of the Andes Amazon Program, September 2025. URL <https://www.maaprogram.org/gold-mining-peru-amazon/>. Published 29 September 2025; observations through mid-2025.
- 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>.
- Robert S. Pindyck. Irreversibilities and the timing of environmental policy. *Resource and Energy Economics*, 22(3):233–259, 2000.
- Mathias Reynaert, Eduardo Souza-Rodrigues, and Arthur A van Benthem. The environmental impacts of protected area policy. *Regional Science and Urban Economics*, 107:103968, 2024.
- Alvin E Roth. Repugnance as a constraint on markets. *Journal of Economic perspectives*, 21(3):37–58, 2007.
- M. Ruiz-Vásquez, P. A. Arias, J. A. Martínez, and J. C. Espinoza. Effects of Amazon basin deforestation on regional atmospheric circulation and water vapor transport towards tropical South America. *Climate Dynamics*, 54:4169–4189, 2020. doi: 10.1007/s00382-020-05223-4.
- Carla Samon Ros. Most of Peru’s gold now comes from booming illegal mining sector. Bloomberg News, republished by MINING.COM, 2 October 2026, October 2026. URL <https://www.mining.com/web/most-of-perus-gold-now-comes-from-booming-illegal-mining-sector/>. Reports IPE estimates: illegally sourced gold was 54% of Peru’s gold export volume, January–June 2026.
- Raúl Sánchez de la Sierra. On the origins of the state: Stationary bandits and taxation in eastern congo. *Journal of Political Economy*, 128(1):32–74, 2020. doi: 10.1086/703989.
- Leonardo A. Saravia, Santiago R. Doyle, and Ben Bond-Lamberty. Power laws and critical fragmentation in global forests. *Scientific Reports*, 8:17766, 2018. doi: 10.1038/s41598-018-36120-w.

- Joseph S. Shapiro. The environmental bias of trade policy. *Quarterly Journal of Economics*, 136(2):831–886, 2021.
- Jean Tirole. Hierarchies and bureaucracies: On the role of collusion in organizations. *Journal of Law, Economics, and Organization*, 2(2):181–214, 1986. doi: 10.1093/oxfordjournals.jleo.a036907.
- Robert M. Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, 21(2):265–293, 1979. doi: 10.1016/0022-0531(79)90031-0. URL <https://www.robertmtownsend.net/content/optimal-contracts-and-competitive-markets-costly-state-verification>.
- United Nations Environment Programme. Global mercury assessment 2018: Key findings. Technical report, United Nations Environment Programme, 2019. URL <https://wedocs.unep.org/handle/20.500.11822/29830>. Global anthropogenic atmospheric emissions inventory for 2015.
- United Nations Group of Experts on the Democratic Republic of the Congo. Letter dated 27 December 2024 from the Group of Experts on the Democratic Republic of the Congo addressed to the President of the Security Council. United Nations Security Council, 2024.
- U.S. Securities and Exchange Commission. Conflict minerals. Final Rule, Release 34-67716, U.S. Securities and Exchange Commission, August 2012. Adopted 22 August 2012. Economic analysis estimates initial compliance costs of \$3–4 billion and ongoing costs of \$207–609 million per year, for about 6,000 issuers.
- Vlado Vivoda, Ron Matthews, and Neil McGregor. A critical minerals perspective on the emergence of geopolitical trade blocs. *Resources Policy*, 89:104587, 2024. doi: 10.1016/j.resourpol.2023.104587. Tables 1–2: refining shares far exceed reserve shares (China 87% REE refining vs 33.9% reserves; 76% cobalt refining vs 1.6% reserves; DR Congo 67% cobalt mining).
- World Gold Council. Gold market primer: Market size and structure, 2026. URL <https://www.gold.org/goldhub/research/market-primer/gold-market-primer-market-size-and-structure>. 18 August; accessed 8 September 2026.