For a Fistful of Gold: Online Appendix

Elies Lucio

22 October 2026

This appendix preserves the supporting material for Sections 2–4 and supplies the new asymmetric-growth and mixed-site proofs for Sections 5–6. The frozen earlier appendix retains superseded model comparisons and their original evidence. Analytical proof, Lean compilation, independent review and manuscript adoption are separate statuses; the accompanying result map records their scope.

Contents

Notation **13**

I Sections 2–3: regional implementation and selective regulation **19**

A Supporting proofs and technical definitions **19**

A.1 Regional implementation: full derivations 19

A.1.1 Regional access, household choice and participation 19

A.1.2 Responses on a linked equilibrium branch 20

A.1.3 Attainment and the terms supporting an allocation 20

A.1.4 Economic bounds for the smooth linked class 22

A.1.5 A unique linked price with trade-dependent survival 23

A.1.6 Effective terms and nominal multiplicity 25

A.1.7 An implemented homogeneous output-cost branch 25

A.1.8 The no-mining corner of the distress benchmark 27

A.2 Resilience: primitives, boundaries and analytical examples 28

A.2.1 Downstream value in the regional output problem 28

A.2.2 Spatial arbitrage and genuine reservation orders 28

A.2.3 A complete trading game and its efficiency restriction 29
A.2.4 Reduction and complete cash accounting 31
A.2.5 The efficient payment bound with the actual pivot basis 32
A.2.6 Attainment against auction and resale deviations 33
A.2.7 Institution choice and the genuine-announcement corollary 34
A.2.8 Private seller information and failure of the broad claim 35
A.2.9 Offered capacity, procurement and selective compliance 36
A.2.10 Primitive operating values and joint retention 37
A.2.11 Individual opportunities and total seller value 38
A.2.12 The original package-depth condition with buyer exclusion 39
A.2.13 Floor conditions and maximum policy effects 40
A.3 Proof of Proposition A.2: Household choice and competitive production . . 45
A.4 Proof of Proposition A.3: Minimum coercion 46
A.5 Proof of Proposition 2.1: Incentives within the committed mechanism . . . 46
A.6 Proof of Lemma 2.5: Whole-buyer critical payments 46
A.7 Proof of Proposition 2.6: Reallocation toward unaffected buyers 46
A.8 Proof of Proposition 2.8: A zero direct laundering response 47
A.9 Proof of Theorem 2.10: The exact exposure identity 47
A.10 Proof of Theorem A.20: Floor conditions for unchanged physical production 47
A.11 Proof of Corollary A.21: A protected marginal return and falling rents . . 48
A.12 Proof of Corollary A.19: The retained-value share 48

B Worked examples and numerical material **49**
B.1 Worked payment comparison 49
B.2 Procurement and unchanged-production example 50
B.3 A positive spread need not cover procurement 51
B.4 Liquidity and the retention reserve 51
B.4.1 Why liquidity alone does not select the reserve 51
B.5 Who buys the gold and who sets its payment 52
B.5.1 Who buys the gold and who determines its payment 52
B.6 Analytical transmission and procurement examples 53
B.6.1 A moving allocation threshold under uncertainty 53
B.6.2 Protected regimes under a fixed type law 54

B.6.3 An exact procurement switch with profitable placement 55
B.6.4 A physical-output switch inside one package band 55
B.6.5 Full support with a protected expected marginal value 56
B.6.6 Strictly partial local output transmission 57
B.6.7 Why a receipt bound does not bound the output derivative 57

II Section 4: certification architecture 60

C Supporting proofs and technical definitions 60

C.1 Certification: general reference account and worked example 61
C.1.1 Physical origin and persistent claims 61
C.1.2 Certification withdrawal without deleting history 62
C.1.3 The horizon of collectible exposure 64
C.1.4 Why a purchaser investigates and reports 64
C.1.5 Independent evidence and the organization of scrutiny 66
C.1.6 Finite diversion opportunities and distributed discovery 67
C.1.7 The funded cost of admission and suppression 68
C.1.8 Entry and the limits of suppression 69
C.1.9 Admission, retrospective removal, and surviving leakage 71
C.1.10 Who pays for credibility 72
C.1.11 A funded market with endogenous scrutiny and fraud 73
C.2 Proof of Proposition C.3: Conditional survival of an admitted fraud 78
C.3 Proof of Proposition C.5: A funded monotonicity condition 78
C.4 Proof of Proposition C.6: Homogeneous channels and saturation 78
C.5 Proof of Proposition C.7: The removal effect 78
C.6 Proof of Proposition C.8: Cost of prevention and a scale limit 78

D Worked examples and numerical material 80

D.1 Mixed-lot withdrawal example 80
D.2 Saved gold calibration 80
D.3 Whole-lot withdrawal and the surviving denominator 81
D.4 Lower leakage and higher funded cost 81

III Sections 5–6: markets, research and government **83**

E Current asymmetric growth and mixed-site policy: proofs **83**

E.1 Result preservation, model comparisons and proof status 83

E.2 Production, actual receipts and aggregation: current proofs 87

E.3 Certification prices: current proofs 90

E.4 Endogenous income, actual dirty rewards and growth: current proofs . . . 92

E.5 Nonempty application and essential boundaries 94

E.6 Asymmetric innovation: ownership, values and research 96

E.6.1 Exclusive assignment, survival and retained vintage quality 96

E.6.2 Proof of surviving-vintage valuation 97

E.6.3 The one-period dirty advantage and later claims 98

E.7 Fixed-pool allocation and the exact support threshold 98

E.7.1 Conditional existence and competitive implementation 98

E.7.2 Proof of the support threshold 99

E.7.3 Verification and absence of additional scientists 100

E.8 Reward divergence, path dependence and sustainable withdrawal 100

E.8.1 Proof of the actual reward envelope 100

E.8.2 Proof of the invariant-region theorem 101

E.8.3 The explicit homogeneous cutoff and finite support 102

E.8.4 Stochastic hitting and its financing limitation 103

E.8.5 Continuing support and unsuccessful withdrawal 103

E.8.6 Dated support conservation 104

E.8.7 Formal reuse and unresolved adapters 104

E.9 Mixed-model fiscal policy, finance and intervention 105

E.9.1 Paid catch-up and the normalized cohort value 105

E.10 Proof of Theorem J.1 108

E.11 A regular connected policy-equilibrium theorem 109

E.12 Proof of Proposition J.2 111

E.13 Preserved certification calibration 112

G Technical remarks on Section 3 of the main text **114**

G.1 Market structure of other minerals 114

G.2 Common-prior reading of the type law . . . 114
G.3 Information measurability of procurement . . . 114
G.4 Single counting of continuation value . . . 114
G.5 Scope of the equilibrium claim and joint deviations . . . 114
G.6 Collusion incentive among coordinated buyers . . . 114
G.7 Extension to procurement and capacity choices . . . 115
G.8 Placement margins under state-dependent sales . . . 115
G.9 Acquisition margin versus production profit . . . 115
G.10 Length of the protected policy interval . . . 115
G.11 Preserved set of cost-minimizing implementations . . . 115
G.12 Laundering share of retained receipts . . . 115

H Technical remarks on Section 4 of the main text 116

H.1 Launch provision for existing stocks . . . 116
H.2 Institutional roles . . . 116
H.3 Record permanence and financial exposure . . . 116
H.4 Content of purchaser losses . . . 116
H.5 Insurance and recourse . . . 116
H.6 Generality of the finite-diversion comparison . . . 116
H.7 Private information and sequential bargaining . . . 116
H.8 Shared collateral and combined roles . . . 117
H.9 Counting at first discovery . . . 117
H.10 Reviewer numbers, common failures and information sharing . . . 117
H.11 Sufficient conditions for the projection property . . . 117
H.12 Fixed concealment targets . . . 117
H.13 Deposit exclusivity, displaced procurement and seizures . . . 117
H.14 Completeness of modeled rewards and liabilities . . . 117
H.15 Scope of withdrawal . . . 118
H.16 Settlement events and economic dates . . . 118
H.17 Definition of contamination ratios . . . 118
H.18 Omitted origins . . . 118
H.19 Participation and solvency of financial support . . . 118
H.20 Historical exposure capacity . . . 118

H.21 Transfers in the resource account . . . 118
H.22 Illustrative parameters . . . 119

IV Complete statements of the current asymmetric model 120

I Market fragmentation and industry change 120

I.1 World trading opportunities . . . 120
I.2 Actual selling receipts and an immobile productive site . . . 121
I.3 Company networks, common wages and income . . . 123
I.4 The certification-date price shock survives with mobile labor . . . 124
I.5 Endogenous growth prices and relative profit divergence . . . 126
I.6 Ownership, surviving certified patents and temporary dirty advantages . . . 128
I.7 One fixed scientist pool and competitive research . . . 130
I.8 Subsidies, reward divergence and a self-sustained green path . . . 131
I.8.1 The funded support threshold at actual values . . . 131
I.8.2 What profitability divergence does and does not establish . . . 132
I.8.3 An invariant private-green region . . . 132
I.9 Support, verification and dated funding . . . 135

J Government action and the transition 135

J.1 The conservation game and its accounts . . . 136
J.2 Optimal taxation subject to private opening and dated finance . . . 138
J.3 The intervention threshold with waiting and chosen taxes . . . 140
J.4 Market interaction and a connected equilibrium class . . . 141
J.5 Safe sites, individual success and cumulative disturbance . . . 142
J.6 A matched comparison of pooling and certification . . . 144

K Further results under the current model 148

K.1 Certification, spending and matched operating rewards . . . 148
K.1.1 The existing mixed-site market . . . 148
K.1.2 A spending identity for the current ownership model . . . 149
K.1.3 Price and operating thresholds are distinct . . . 152
K.1.4 Closed forms within the current specification . . . 153
K.2 Complete rewards and the preservation of private research . . . 154

K.2.1 Actual effort, values and support . . . 155
K.2.2 Exact preservation and cost headroom . . . 156
K.2.3 The existing reversal examples . . . 157
K.3 Sharper path dependence within the existing model . . . 158
K.3.1 Existing hypotheses and the exact marginal comparison . . . 159
K.3.2 A tighter sufficient cutoff and an explicit grant bound . . . 160
K.3.3 Withdrawal dates and stochastic hitting . . . 162
K.3.4 The exact limiting trichotomy and the boundary case . . . 163
K.3.5 Using surviving rents without changing patent ownership . . . 164
K.4 From the price shock to mine income and protection payments . . . 166
K.4.1 An exact link for the coerced operation . . . 166
K.4.2 What can be concluded about protection income . . . 167
K.5 Monotonicity beyond the installed-input specialization . . . 168
K.6 An exact intervention threshold with tax-dependent success . . . 169

V Preserved 29 October comparison model 171

L Market fragmentation and industry change 171

L.1 World trading opportunities . . . 171
L.2 North and South as buying markets . . . 172
L.3 Regional supply and the income it creates . . . 173
L.4 The North–South economy: pooling and certification . . . 174
L.4.1 The price shock after certification . . . 175
L.4.2 Segmentation and the global policy ceiling . . . 176
L.4.3 Market size and the return to clean capacity . . . 176
L.5 Certified and outside prices . . . 177
L.6 Whole-buyer selling and processing access . . . 179
L.7 Price feedback and the opportunity to invest . . . 180
L.8 Machines, production and the return to an improvement . . . 181
L.9 From production to the investor's innovation value . . . 182
L.10 Schumpeterian innovation and the direction of research . . . 183
L.11 Path dependence, subsidy thresholds and private takeover . . . 185
L.11.1 The payment needed at market values . . . 185

L.11.2 An explicit inherited-capacity threshold 186
L.12 Support, verification and funding 188
L.13 From the available standard to government action 189

M Government action and the transition **191**
M.1 Tax revenue and protection income at the same prices 191
M.2 Solving the intervention comparison 193
M.2.1 A constant-flow benchmark 193
M.2.2 The threshold on the changing market path 194
M.3 Market feedback and the transition 195
M.4 Funded livelihoods and protected sites 196

N Supporting proofs and technical definitions **198**
N.1 Exact geographic aggregation and household reconstruction 198
N.1.1 Measure, scale and the cost infimum 198
N.1.2 Contracts, accounting moments and invariance 200
N.1.3 Measurable and integrable return to households 201
N.1.4 Entry, finite coverage and subsistence restrictions 204
N.2 Common-price equilibrium, composition and four accounts 207
N.2.1 A complete timing and ownership specification 207
N.2.2 Fraud and diligence from individual choices 208
N.2.3 Global price solution and regularity 209
N.2.4 Four linked accounts 209
N.2.5 Scaling, owners and benchmark recovery 210
N.3 Proofs for the active nonlinear world 211
N.3.1 The cell, its regional cost, and its true physical account 211
N.3.2 A continuum menu cell with active resale and positive retained delivery 213
N.3.3 A separate finite whole-buyer VCG construction 215
N.3.4 The equilibrium operator, reconstruction and scale 217
N.3.5 Proof of the explicit certification connection 218
N.3.6 A uniform connected stage for an infinite growth path 219
N.3.7 A financed faithful retail provenance subclass 222
N.3.8 Commitment, limits and counterexamples 225

N.4 Paid progress, research ownership, and global institutional dynamics 226
N.4.1 Dates, ownership, scale, and the stage-market interface 226
N.4.2 A complete verification workload and its financing 229
N.4.3 Actual research prizes, including surviving private liabilities 231
N.4.4 Reclearing markets with a positive mass of governments 233
N.4.5 Scarcity can strengthen protection: counterexample and repair . . . 236
N.4.6 When intervention becomes attractive 238
N.4.7 Returning the theorem to households and a physical forest 238
N.4.8 Household income behind aggregate supply 241
N.5 Proof of Proposition N.26: Exact geographic aggregation 242
N.6 Proof of Theorem N.27: Joint common-price equilibrium 242
N.7 Proof of Proposition N.28: An exact active nonlinear class 242
N.8 Retained baseline market definitions 242
N.9 Regional production and world supply 243
N.9.1 Who produces and who receives the rents 244
N.9.2 Household income behind aggregate supply 245
N.10 Certification and world-market equilibrium 246
N.10.1 The contract buyers purchase 246
N.10.2 The incentives to investigate and report 247
N.10.3 Market composition and the destination of production 247
N.10.4 Prices, production and admission 248
N.10.5 Production rents and surviving contamination 249
N.11 Retail market power and the world price 250
N.11.1 Access to processing and the return to production 251
N.11.2 Substitution across the certified, outside and retail markets 253
N.11.3 The income available to innovators and governments 255
N.12 Technical change and the incentive to govern 256
N.12.1 Maintaining credibility as production expands 257
N.12.2 Private returns and the direction of research 257
N.12.3 Growth, market access and resources 259
N.12.4 Government choices and their market consequences 260
N.12.5 Reform and the rents of remaining illegal producers 261
N.12.6 The decision to establish control 261

N.12.7 Household livelihoods and protected sites 262

O Worked examples and numerical material **265**

O.1 Private and supported research scenarios 265

O.1.1 Private and supported research in two scenarios 265

O.2 Joint dynamic numerical comparisons 265

O.2.1 Joint changes in candidate extraction 265

O.3 Explicit finite whole-buyer cell 266

O.4 An equilibrium and its physical accounts 267

O.5 Aggregation and the distribution of market power 267

O.6 Supply and profit use different moments 268

O.7 A physical damage account beyond the supply moment 268

O.8 Entry costs and the power-envelope comparison 268

O.9 A minimum operating output and participation 268

O.10 Checking the interior common-price equilibrium 269

O.11 Scarcity can raise a remaining producer's rent 269

O.12 Waiting before intervention becomes attractive 269

O.13 Exact common-price equilibrium construction 269

O.13.1 An exact equilibrium and arbitrarily small positive leakage 269

O.14 Observed evidence and saved numerical comparisons 271

O.14.1 Observed outcomes and construction 271

O.14.2 Certification: rates, physical masses and saved sensitivity 272

O.14.3 Dynamic scenarios, availability and finance 273

O.14.4 Reproduction and statement correspondence 275

P Core economic proofs **276**

P.1 Subsistence income and effort 277

P.1.1 The subsistence labor response 277

P.1.2 Consolidated resources 277

P.1.3 The subsistence financing floor 277

P.2 Different prices, the same output and profit 277

P.2.1 Attained effective terms, nonunique gross quotes 277

P.2.2 Regional output as an implemented choice 278

P.3 Buffers against selective regulation . . . 278
P.3.1 The opportunity floor . . . 278
P.3.2 The remaining-market floor . . . 279
P.3.3 Global output resilience . . . 279
P.3.4 Protection by unaffected payment-setting orders . . . 279
P.3.5 Falling rents with unchanged extraction . . . 279
P.4 Partial transmission . . . 280
P.4.1 Local production transmission . . . 280
P.4.2 An analytical partial-transmission condition . . . 280
P.5 Ceilings on receipt, profit and production losses . . . 280
P.5.1 The exact floor and maximum loss of receipts . . . 280
P.5.2 Maximum profit loss and a quantity bound . . . 281
P.5.3 A resource bound on policy effects . . . 282
P.5.4 Global policy ceiling . . . 283
P.6 Physical and incentive architecture . . . 284
P.6.1 Conservation and retrospective exclusion . . . 284
P.6.2 Investigation and voluntary disclosure . . . 284
P.6.3 Funded capture with optional suppression . . . 284
P.6.4 Positive-trade interior equilibrium . . . 285
P.7 The certification price shock . . . 285
P.8 Patent ownership, research and withdrawal . . . 286
P.9 Taxation, intervention and physical conservation . . . 286

Q Core economic proofs **286**

Q.1 Subsistence income and effort . . . 287
Q.1.1 The subsistence labor response . . . 287
Q.1.2 Consolidated resources . . . 287
Q.1.3 The subsistence financing floor . . . 287
Q.2 Different prices, the same output and profit . . . 287
Q.2.1 Attained effective terms, nonunique gross quotes . . . 287
Q.2.2 Regional output as an implemented choice . . . 288
Q.3 Buffers against selective regulation . . . 288
Q.3.1 The opportunity floor . . . 288

Q.3.2 The remaining-market floor . . . 289
Q.3.3 Global output resilience . . . 289
Q.3.4 Protection by unaffected payment-setting orders . . . 289
Q.3.5 Falling rents with unchanged extraction . . . 289
Q.4 Partial transmission . . . 290
Q.4.1 Local production transmission . . . 290
Q.4.2 An analytical partial-transmission condition . . . 290
Q.5 Ceilings on receipt, profit and production losses . . . 290
Q.5.1 The exact floor and maximum loss of receipts . . . 290
Q.5.2 Maximum profit loss and a quantity bound . . . 291
Q.5.3 A resource bound on policy effects . . . 292
Q.5.4 Global policy ceiling . . . 293
Q.6 Physical and incentive architecture . . . 294
Q.6.1 Conservation and retrospective exclusion . . . 294
Q.6.2 Investigation and voluntary disclosure . . . 294
Q.6.3 Funded capture with optional suppression . . . 294
Q.6.4 Positive-trade interior equilibrium . . . 295
Q.7 The certification price shock . . . 295
Q.7.1 The faithful North–South price comparison . . . 295
Q.8 Research support . . . 296
Q.8.1 Paid verification and clean-directed mineral research . . . 296
Q.8.2 Actual-value and uniform research-support thresholds . . . 297
Q.9 Path dependence and private takeover . . . 298
Q.9.1 Capacity, private takeover and the duration of support . . . 298
Q.10 The government intervention threshold . . . 299
Q.10.1 Eventually strict intervention on the recleared path . . . 299
Q.10.2 The fiscal threshold with waiting and repeated attempts . . . 300
Q.11 Conservation . . . 300
Q.11.1 What governance does and does not conserve . . . 300

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	2–5
$\bar{q}$	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	2–5
$\mathcal{C}(q)$, $\Gamma(q)$	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, 5–6
h	Certified quote, the commodity invoice in the certified market†	4, 5–6
ι	Illegal share of certified mineral	4–5, D
Ω	Public market state, including committed continuation arguments	4–5, 7
π	Share of fraudulent admissions; posterior probability of fraud†	4–5
η	Elasticity of regional output with respect to its unit receipt†	5–6
E_C, E_D, E_R	Expenditure on certified, outside and retail purchases	4, 5
X, Y	Certified and outside purchases	5
k	Origin of mineral†	4–5
t	Date	4, 5–6
Subscripts C, D	Certified (clean) and outside (dirty)	throughout

Section 1: 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 1.6)
$\mathcal{A}_r^{\text{safe}}$	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; newly retained cash; saved cash unavailable for current spending
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 2: Selective regulation and the return to extraction

Symbol	Meaning
θ_E	Gross world-market price of physical gold outside the gate, per standardized unit
$\mathcal{R}$, $\mathcal{U}$	Regulated and unaffected buyers†
$\mathcal{I}$, N	Buyers; their number
$H(y)$	Retention value of unsold stock y, through laundering, storage or disposal†
$\bar{k}_i$, ω_i, ν	Public capacity and private type of buyer i; joint 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 kth 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
R_n, W_n	Realized and expected gross stock value with offered capacity n, including buyer payments and retained value
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 $\rho_1(a), \ldots, \rho_n(a)$
$\rho(q)$	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 2.5†
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
χ_n, $\bar{\chi}_n$	Realized and expected exposure of payments to regulation; $\bar{\chi}(\Xi, \tau)$ evaluates exposure at implementation Ξ
α_R	Cumulative receipt loss relative to a charge on all offered units
a in Theorem 2.14	Gradient of regional output with respect to feasible implementation coordinates; distinct from unoffered stock†
$\mathcal{H}$, Λ, γ	Profit Hessian; regional adjustment factor; marginal payment exposure in the local production result

Section 3: 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
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^{sil}, c^{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
ι^{in}, ι^{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

Current Sections 5–6: asymmetric growth and the mixed site

Symbol	Meaning
β_m, α_H, α_L	Machine-composite exponent and paid headquarters/local-labor exponents; $\alpha_\Sigma = \alpha_H + \alpha_L < 1$
H_r^{site}, $H_{f,t}$, $L_{fr,t}$	Fixed productive site, shared paid HQ input, mobile ordinary branch labor
w_t^L, $c_t w_t^R$	Endogenous ordinary wage and scientist wage

Symbol	Meaning
$A_{C,t}$, $A_{D,t}$, a_v	Mean sectoral line qualities and a particular surviving patent vintage
$x_{r,a}$, χ_r	Paid catch-up at age since opening and site technology factor
$x^M(m)$, q_j^M	Machine use and machine quote; distinct from catch-up $x_{r,a}$
B_f^F, C_{fr}^F, Ψ	Company coefficient, branch coefficient, optimized corporate operating value
$\mathcal{J}_j$, $\mathcal{K}_j$	Company-network moments for labor and gross production; not a regional mean
$\rho_{C,t}$, δ_R	Research-implied displacement probability and research-owner discount
$J_{i,t}$, $\widehat{J}_{i,t}$	Conditional success-minus-failure private claim value and its value divided by c_t
$h_i(s)$, $R_i(s)$, N_i	Research success, convex laboratory cost, fixed project-owner mass
S, S_V, $G_{C,t}^{\min}$	Fixed total scientists, paid verification provision, minimum additive research support
τ_r^G, ξ_a	Output royalty and distinct net-profit tax; neither is the Section 3 compliance cost
$z_{r,t}^G$, $z_{r,t}$	Boolean public control and the full physical implemented allocation
$U_{r,e}^{ctrl}$, $V_{r,t}$	Value at control/opening date and value before control with the option to wait
$K_{r,t}^G$, $P_{r,t}$, δ_G	Attempt cost, next-period success probability, government discount
B_t, d_t	Endogenous monetary opportunity scale and $\delta_G B_{t+1}/B_t$
$S_{r,t}^{sub}$, $R_{r,t}^G$	Exact household subsistence bill and net current fiscal flow
$M_{a,t}^D$, $v_{a,t}^G$	Ungoverned regional mass and attempted fraction; not scientist mass
g_D, $\mathcal{B}$	Physical disturbance growth envelope in the conservation result and irreversible-damage allowance

Scope of reused symbols

Symbol	Distinct meanings retained explicitly
τ, τ_r^G, ξ_a	Section 3 buyer compliance charge; Section 6 output royalty; alternative profit tax
h, $h_i(s)$	Certified commodity quote; research success probability
x, $x_{r,a}$, $x^M(m)$	Local-good notation in Section 2; catch-up fraction; machine use
$H(y)$, H_r^{site}, $H_{f,t}$	Retention value in Section 3; fixed productive site; paid shared headquarters
π, π^v	Fraud share in Sections 4 and its supporting proofs; patent-vintage profit
η_v, η	Household leisure weight; supply elasticity on its declared power-cost branch

Symbol	Distinct meanings retained explicitly
a_v, a	Household opening wealth in Section 2 or vintage quality in the local patent discussion; catch-up age when subscripted by entry date
L, L_E, L_L, $L_{fr,t}$	Laundering in Section 3; purchaser losses in Section 4; ordinary branch workers
N_i	Rival-order value in the stated Section 3 proof; research project-owner mass in Section 5
S_t, S, S_V	Live claims in Section 4; fixed scientists; verification use of that pool
z, z^G	Full implemented physical allocation in Section 2 and the safe-set criterion; Boolean control in Section 6
κ_r, c_{ur}	Procurement-cost parameter including household rent; real unit operating-input cost in the installed-input specialization

Part I

Sections 2–3: regional implementation and selective regulation

A Supporting proofs and technical definitions

A.1 Regional implementation: full derivations

A.1.1 Regional access, household choice and participation

The following results provide the regional characterizations used in Section 1. The proofs of Propositions A.2 and A.3 appear in the proof subsections A.3 and A.4, respectively.

Lemma A.1 (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 1.1, 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. □

In the household conditions below, λ_v is the multiplier on the budget constraint and ζ_v is the multiplier on the time constraint.

Proposition A.2 (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*

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

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. See Appendix A.3.

Proposition A.3 (Minimum coercion)**.** *If coercion affects the household only through* (7)*, an implementable allocation uses*

$$e_v^* = \Delta_v^{-1}((\bar{U}_v - u_v)_+) \tag{2}$$

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 with* $\Delta_v'(e_v^*) \neq 0$*,* $de_v^*/ds = -u_{v,s}/\Delta_v'(e_v^*)$ *for a regional control* s*.*

Proof. See Appendix A.4.

A.1.2 Responses on a linked equilibrium branch

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,

$$x_s = -\mathcal{F}_x^{-1}\mathcal{F}_s, \qquad q_s = \sum_v (G_{vL} l_{v,s}^m + G_{vg} g_{v,s}) \tag{3}$$

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.

The gate's objective is defined in Section 1; its derivative accounts for the induced household and trading responses.

On a smooth regime its derivative in a control s_j is

$$\begin{aligned}\Pi_j = (W_q - \theta_B) q_j - \mathbf{1}_{\{s_j = \theta_B\}} q + \mathbf{1}_{\{s_j = \mu_z\}} Z + \mu_z Z_j \\ + \mathbf{1}_{\{s_j = \mu_g\}} I + \mu_g I_j - T_j - \sum_v K_v'(e_v) e_{v,j}\end{aligned} \tag{4}$$

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.

A.1.3 Attainment and the terms supporting an allocation

The following class supplies the sufficient conditions used for the regional cost and output comparison.

Assumption A.4 (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

$$d_A := \sum_{v \in A} \alpha c_v^z c_v - \sum_{v \in A} a_v > 0 \tag{5}$$

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 A.5 (Attained regional cost)**.** *Under Assumption A.4, 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.*

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, follow in the next subsection.

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.

A.1.4 Economic bounds for the smooth linked class

The result concerns a nonempty equilibrium graph under Assumption A.4. 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 A.5. Let $H = M + \sum_v a_v$. The exact resource account implies $\alpha c_v^z C_v, c_v^g g_v, b_v \le H$ individually. Thus C_v, g_v, b_v have explicit upper bounds, and $r_v \le 1$. Put $\underline{U}_v = \bar{U}_v - \Delta_v(\bar{e}_v)$. Participation gives

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

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,

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

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

$$\theta_B G_v = w_v l_v^m/\gamma_v, \qquad p_v^g g_v = \kappa_v \theta_B G_v \tag{6}$$

Hence $\theta_B q \le D := \sum_{A_v^m > 0} \overline{w}_v/\gamma_v$. The external cash account and (5) imply $\theta_B q \ge d_A$. The resource input bounds and $l_v^m \le 1$ give $q \le \bar{q} := \sum_v A_v^m (H/c_v^g)^{\kappa_v}$. Consequently

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

For nonnegative profit, $C(z) \ge d_A$ and $W \le \bar{v}q$ 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 $G_T \ge 1$, $q \le A_T G_T^{\bar{\kappa}}$ and $C(z) \ge c_g^{\min} G_T - a_T$. If additionally $G_T \ge 2a_T/c_g^{\min}$, nonnegative profit implies $(c_g^{\min}/2)G_T \le \bar{v} A_T G_T^{\bar{\kappa}}$. Thus all such choices have

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

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

At least one of the N_m capable miners has $G_v \ge q_0/N_m$, hence $g_v \ge [q_0/(N_m A_v^m)]^{1/\kappa_v}$. Equation (6) 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. □

A.1.5 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. Here cash denotes additions to savings: zero cash means that households retain no current income. Define

$$p_- := \frac{Bc_1}{A_x - a_x c_1}, \qquad p_+ := \frac{D/c_2 - B}{a_x}, \qquad t_1(p) = \frac{pA_x}{a_x p + B}, \quad t_2(p) = \frac{D}{a_x p + B}$$

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

$$b_{01}/\beta_1 > A_x p_+, \qquad b_{02}/\beta_2 > D \tag{7}$$

Theorem A.6 (Unique linked price from a monotone survival equation)**.** *On (p_-, p_+) let $l_i(p) \in (\nu_i/(\nu_i + \eta_i), 1)$ solve*

$$t_i(p) l_i^{\nu_i - 1}[(\nu_i + \eta_i) l_i - \nu_i] = \eta_i c_i \tag{8}$$

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 $\lambda_i \le \beta_i/b_{0i}$, whereas $1 = \lambda_i P(C_i - c_i) < \lambda_i \text{income}_i < 1$ by (7). 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 (8). 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. □

A.1.6 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

$$\begin{aligned} \max \quad & \log(C - c) + \eta \log r + \beta \log(b + b_0) \\ \text{subject to} \quad & dC + l + r = 1, \quad BC + b = wl, \\ & C > c, \quad r > 0, \quad b, l \geq 0 \end{aligned} \tag{9}$$

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

$$C - c = \frac{M}{D(B + wd)}, \qquad r = \frac{\eta M}{Dw}, \qquad b + b_0 = \frac{\beta M}{D} \tag{10}$$

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.7 (Attained effective terms, nonunique gross quotes)**.** *With bounded coercion, finite outside utility and continuous nonnegative transport and coercion costs, every nonempty cost sublevel in (9) is compact in physical allocations and effective terms (w, B). Every nonempty output fiber attains its cost minimum. Its nominal implementation set contains*

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

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

The full proof is in the paper appendix, Section P.2.1.

A.1.7 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

$$q = a(A)g^{\alpha_m}, \qquad 0 < \alpha_m < 1 \tag{11}$$

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.8 (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*

$$\begin{aligned} \mathcal{C}(q; A) &= \frac{c_g}{\alpha_m} a(A)^{-1/\alpha_m} q^{1/\alpha_m}, \\ \theta_B(q; A) &= \frac{c_g}{\alpha_m} a(A)^{-1/\alpha_m} q^{1/\alpha_m - 1} \end{aligned} \tag{12}$$

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

$$a(A) = \left(\frac{c_g}{\kappa \alpha_m^2}\right)^{\alpha_m} A^{1-\alpha_m}, \qquad \kappa > 0 \tag{13}$$

Then the envelope becomes

$$\mathcal{C}(q; A) = \frac{\kappa}{1 + 1/\eta} A^{-1/\eta} q^{1+1/\eta} \tag{14}$$

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

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

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 (12). 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 (13) 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. □

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.1.8 The no-mining corner of the distress benchmark

The interior expression in Proposition 1.3 is used only when mining labor is positive and the household survives. If wealth allows zero mining, the no-mining allocation solves

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

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.

A.2 Resilience: primitives, boundaries and analytical examples

A.2.1 Downstream value in the regional output problem

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

$$W(q) = \max_{n \in \mathcal{N}(q)} W_n(q), \qquad \mathcal{N}(q) \subseteq \{0, \dots, \lfloor q/\bar{q} \rfloor\} \tag{15}$$

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.

The regional choice is characterized in Theorem 1.5.

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.

A.2.2 Spatial arbitrage and genuine reservation orders

The results in this subsection and those that follow prove the conditional institutional comparison in Theorem 2.3. They do not replace the broader, committed-institution resilience results.

Lemma A.9 (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*

$$H(y + \bar{q}) - H(y) = \bar{q}\theta_E - [V(y + \bar{q}) - V(y)]. \tag{16}$$

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 (21), including residual stock and channel capacities.

Proof. The feasible alternatives to selling the additional package are the joint plans in (16). Their optimized value difference is therefore the stated increment. Substituting the two pure-laundering endpoint values gives (16); 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. □

The worked example is in Online Appendix B.3.

A.2.3 A complete trading game and its efficiency restriction

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

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

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

$$F_H(k; b) = \sum_i v_i(k_i; b_i, 0) + H(q - \bar{q} \textstyle\sum_i k_i), \qquad S_H(b) = \max_{k \in \mathcal{K}_q} F_H(k; b). \tag{17}$$

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 $(\delta_0, \delta_1, \ldots, \delta_N)$ 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 A.10 (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 A.10.

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 involving 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.

A.2.4 Reduction and complete cash accounting

Lemma A.11 (Direct representation of the complete equilibrium)**.** *Every outcome in Definition A.10 induces an efficient, buyer Bayesian incentive-compatible and interim individually rational direct mechanism. Its expected seller continuation value is*

$$\mathbb{E}[T_0^{\text{net}} + H(y^f)] = \mathbb{E}S_H(b) - \sum_i \mathbb{E}U_i(b_i), \tag{18}$$

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 A.10.

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

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

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 (18). 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. □

A.2.5 The efficient payment bound with the actual pivot basis

Lemma A.12 (Zero is a most-reluctant type)**.** *For every b_{-i},*

$$S_H(0, b_{-i}) = S_{H,-i}(b_{-i}). \tag{19}$$

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 (19). Substitution of the actual pivot payment in (18) 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. □

Theorem A.13 (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*

$$U_i^M(b_i) = U_i^V(b_i) + c_i, \tag{20}$$

$$\mathbb{E}R^M = W_Q(q, 0) - \sum_i c_i \leq W_Q(q, 0). \tag{21}$$

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

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

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 A.12 implies $g(0) = \mathbb{E}_{-i} S_{H,-i}$ and $g(1) - g(0) = U_i^V(t)$. The type t was arbitrary, proving (20); 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 (21). This proves optimality of the exclusion basis here without adding participation fees or replacing the mechanism by another generalized VCG. □

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.

A.2.6 Attainment against auction and resale deviations

Lemma A.14 (Complete-game attainment)**.** *For the genuine-H, all-offered pivot auction followed by the terminal consent protocol in Section A.2.3, 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

$$p_i^a + H(y_a) + \sum_{j \neq i} v_j(a_j; b_j, 0) = S_{H,-i}(b_{-i}). \tag{22}$$

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

$$v_i(a_i; b_i, 0) - p_i^a + s_i \leq S_H(b) - S_{H,-i}(b_{-i}), \tag{23}$$

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 (22)–(23) 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 (23), 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. □

A.2.7 Institution choice and the genuine-announcement corollary

Proof of Theorem 2.3. Lemma A.11 represents any admissible efficient complete outcome as an efficient BIC and IR direct mechanism with the exact seller account. Theorem A.13 bounds its expected continuation payoff by $W_Q(q, 0)$. Lemma A.14 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. □

Proof of Corollary 2.4. 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 2.3. At the genuine announcement select the attaining equilibrium in Lemma A.14. 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. □

A.2.8 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 A.15 (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 A.13, imply

$$U_B'(b) = \Pr(c \leq b) = b, \qquad U_S'(c) = -\Pr(b \geq c) = -(1 - c)$$

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

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

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. □

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.

A.2.9 Offered capacity, procurement and selective compliance

Proposition A.16 (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 A.10.*

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 semicontinuity 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 (19) 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. □

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.

The worked example is in Online Appendix B.4.

A.2.10 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 (16) 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.

A.2.11 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

$$R_n - H(q) = \sum_i [p_i - D_i(y)] + \left\{ \sum_i D_i(y) - [H(q) - H(y)] \right\} \tag{24}$$

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 $W_n(q, \tau) - \mathcal{C}(q)$.

A.2.12 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

$$K = (u-1)a + r, \qquad a = \lfloor K/(u-1) \rfloor, \qquad 0 \leq r < u-1 \tag{25}$$

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

$$\beta_{\mathcal{U}} = \begin{cases} \min_{h \in \mathcal{U}} b^a_h, & r = 0, \\ b^1_{\mathcal{U},[r+1]}, & a = 0, \\ \min\{\min_{h \in \mathcal{U}} b^a_h, b^{a+1}_{\mathcal{U},[r+1]}\}, & a \geq 1, r > 0. \end{cases} \tag{26}$$

Proposition A.17 (A sufficient unaffected-demand depth)**.** *For any fixed information law and acquisition quote,*

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

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 $\min_h b^a_h$, 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. □

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.

The complete analytical construction is in Online Appendix B.6.

A.2.13 Floor conditions and maximum policy effects

This subsection supplies the global comparisons used in Section 2.8. Buyer values, the type law, retention, regional implementation opportunities and permitted capacities are fixed. All quantity comparisons concern physical stock, including its fractional residual, and all profit comparisons include exit at its specified payoff.

The exact surviving-market floor. For fixed q, n, ω, remove every regulated buyer and run the same whole-buyer pivot auction with unaffected buyers, genuine retention H and the same offered capacity. Write its seller value as $R_n^{\mathcal{U}}(q,\omega)$ and define

$$\underline{W}_n(q) = \mathbb{E}_\nu R_n^{\mathcal{U}}(q,\omega), \qquad M_n(q) = W_n(q,0) - \underline{W}_n(q). \tag{27}$$

This is an auction with the surviving owners, not an auction that reassigns their orders to independent fictitious bidders. It preserves whole-buyer exclusion and the joint use of retention channels. It is therefore a sharper floor than retaining all stock alone.

Proposition A.18 (The exact floor and maximum loss of receipts)**.** *For every feasible fixed stock and offered capacity,*

$$\begin{aligned} H(q) \le \underline{W}_n(q) \le W_n(q,\tau) \le W_n(q,0), \\ 0 \le W_n(q,0) - W_n(q,\tau) \le \min\{n\bar{q}\tau, M_n(q)\}. \end{aligned} \tag{28}$$

Moreover,

$$\lim_{\tau\to\infty} W_n(q,\tau) = \underline{W}_n(q), \qquad \sup_{\tau\ge 0}\{W_n(q,0) - W_n(q,\tau)\} = M_n(q). \tag{29}$$

Thus $M_n(q)$ is the exact maximum loss of expected receipts, understood as a supremum when the type distribution has unbounded support.

For $n \ge 1$, let $B_{\mathcal{R}}(\omega) = \max_{i\in\mathcal{R}}(b_i^1)_+$, with value zero when $\mathcal{R}$ is empty. A sufficient statewise saturation charge is

$$T_n(q,\omega) = \frac{[B_{\mathcal{R}}(\omega) - e_{\mathcal{U},H}(q,n,\omega)]_+}{\bar{q}}. \tag{30}$$

For every $\tau \ge T_n(q,\omega)$, seller value equals $R_n^{\mathcal{U}}(q,\omega)$. Set $T_0 = 0$. Consequently the expected floor is attained for all charges at least $\operatorname{ess\,sup}_\nu T_n(q,\omega)$ whenever this bound is finite. The simpler bound

$$\overline{T} = \frac{\operatorname{ess\,sup}_\nu B_{\mathcal{R}}(\omega)}{\bar{q}} < \infty \tag{31}$$

is sufficient simultaneously for every stock and every allowed capacity.

Proof. See Appendix P.5.1.

Corollary A.19 (The retained-value share under selective regulation)**.** *Fix physical stock q, offered capacity n, the type law and the priority at ties. Let $y_n(q,\omega,\tau)$ be the stock*

retained in the efficient auction allocation. It is weakly increasing in τ state by state. If $H \geq 0$ on the relevant retained stocks, then on every interval with $W_n(q,\tau) > 0$ the share

$$s_H(q,n,\tau) = \frac{\mathbb{E}_\nu H(y_n(q,\omega,\tau))}{W_n(q,\tau)} \tag{32}$$

is weakly increasing in τ. When laundering is the sole valued retained use, this is the share of receipts supplied by net laundering proceeds. The statement fixes stock and capacity and asserts neither a strict increase nor monotonicity after their reoptimization.

Proof. See Appendix A.12.

The cutoff checks the full allocation and every buyer-exclusion allocation. It may therefore be later than the charge at which the last regulated buyer stops receiving gold: an unsuccessful regulated order can still determine somebody else's payment. Conversely, a strong unaffected book can make the floor bind before regulated values reach zero. The cutoff is sufficient, not necessarily the smallest charge at which receipts stop falling. Equation (29) gives the exact ceiling even when no finite uniform cutoff exists.

The floor also has transparent conservative certificates. Retention alone gives $\underline{W}_n(q) \geq H(q)$. For $n \geq 1$, if, almost surely, $e_{\mathcal{U}}(n,\omega) > \rho_1(q - n\bar{q})$, unaffected demand alone fills the entire offered capacity even when any one buyer is removed. All n packages then sell externally in the surviving auction, and the remaining-market bound gives

$$\underline{W}_n(q) \geq H(q - n\bar{q}) + n\mathbb{E}_\nu e_{\mathcal{U},H}(q,n,\omega). \tag{33}$$

The premise compares unaffected demand with the highest optional retention order. Without a guaranteed number of sales, multiplying a per-package floor by the offered count would overstate the protected receipts. Positive protected receipts net of implementation cost can rule out exit; they do not alone select a particular positive output.

An explicit floor test for unchanged extraction. Allow the gate to reoptimize offered capacity and write

$$W(q,\tau) = \max_{n \in \mathcal{N}(q)} W_n(q,\tau), \qquad \underline{W}(q) = \max_{n \in \mathcal{N}(q)} \underline{W}_n(q). \tag{34}$$

Since $\mathcal{N}(q)$ is finite at every fixed stock, these maxima commute with the decreasing policy limit. The exact maximum loss with capacity reoptimized at that stock is therefore $W(q,0) - \underline{W}(q)$, including under unbounded type support. For active quantities define $V(q,\tau) = W(q,\tau) - \mathcal{C}(q)$ and $\underline{V}(q) = \underline{W}(q) - \mathcal{C}(q)$. Include exit explicitly by setting $V(0,\tau) = \underline{V}(0) = \Pi^{\text{exit}}$ at its specified, policy-independent payoff. Let $q_0 > 0$ be a baseline maximizer, let n_0 attain $W(q_0,0)$, and define

$$\begin{aligned} \Delta_0(q) &= V(q_0,0) - V(q,0), \\ D_0 &= W(q_0,0) - \underline{W}(q_0), \qquad B_0(\tau) = \min\{n_0\bar{q}\tau, D_0\}. \end{aligned} \tag{35}$$

The original output can lose at most $B_0(\tau)$ after optimizing its selling decision. Every alternative's profit weakly decreases as well.

Theorem A.20 (Floor conditions for unchanged physical production)**.** *At any policy τ, it is sufficient for unchanged optimal output that*

$$\Delta_0(q) \geq B_0(\tau) \quad \textit{for every feasible } q \neq q_0, \quad \textit{including exit.} \tag{36}$$

Strict inequalities make q_0 the unique optimal physical output. In particular, the policy-independent floor test

$$\underline{W}(q_0) - \mathcal{C}(q_0) \geq V(q, 0) \qquad \textit{for every } q \neq q_0 \tag{37}$$

preserves q_0 for every charge. Again, strict inequalities against all different quantities give uniqueness.

If the baseline already receives its surviving-market floor,

$$W(q_0, 0) = \underline{W}(q_0), \tag{38}$$

then q_0 remains optimal for every $\tau \geq 0$ without further restrictions on the other quantities. A unique baseline quantity remains unique. More generally, if all baseline maximizers have quantity q_0 and one baseline selling choice attains its floor, no different output can become optimal.

Proof. See Appendix A.10.

This is the sense in which a binding receipt floor prevents a production response. The level $H(q_0)$ or an unaffected payment floor being high is insufficient; its protected profit must cover the competing quantities, or it must already bind at a baseline optimum. A floor above the procurement quote only establishes a placement margin, whereas the test subtracts the full regional cost and compares every feasible output. Replacing $\underline{W}(q_0)$ in (37) by a verified lower bound from retention or (33) gives a more conservative sufficient test.

If $\delta = \inf_{q \neq q_0} \Delta_0(q) > 0$, a convenient strict certificate is $B_0(\tau) < \delta$: it holds for every charge if $D_0 < \delta$, and otherwise for $\tau < \delta/(n_0 \bar{q})$ when $n_0 > 0$. With continuous physical output, however, nearby quantities commonly have arbitrarily small baseline gaps, so $\delta = 0$. The finite-gap certificate must then be replaced by a binding floor or by a comparison of differential losses, as in Theorem 2.11. An unchanged integer package count does not establish unchanged physical output.

A marginal retention floor can preserve output while rents fall. There is a useful explicit continuous-output case. Suppose $H(y) = hy$ on all stocks under consideration,

with $h \geq 0$. For a fixed offered count n, every optional seller order equals $h\bar{q}$, independently of physical stock. Consequently

$$W_n(q, \tau) = hq + \psi_n(\tau), \qquad \psi_n(\tau) \geq 0, \tag{39}$$

where ψ_n depends on the buyer population, capacity and policy but not on q. For each permitted capacity branch $I_n = \{q > 0 : n \in \mathcal{N}(q)\}$, suppose

$$J_n = \max_{q \in I_n} \{hq - \mathcal{C}(q)\}, \qquad q_n \in \arg\max_{q \in I_n} \{hq - \mathcal{C}(q)\} \tag{40}$$

are attained. Let $\underline{\psi}_n = \lim_{\tau \to \infty} \psi_n(\tau)$. An interior differentiable strict optimum satisfies $\mathcal{C}'(q_n) = h$; on an interval branch, strict convexity makes that branch optimum unique.

Corollary A.21 (A protected marginal return and falling rents)**.** *Under these affine-retention conditions, a baseline branch n_0 and its output q_{n_0} remain optimal at every policy if*

$$J_{n_0} + \underline{\psi}_{n_0} \geq \max \left\{ \Pi^{\text{exit}}, \sup_{n \neq n_0} [J_n + \psi_n(0)] \right\}. \tag{41}$$

If the branch optimum is unique and all displayed comparisons are strict, physical output is uniquely q_{n_0} for every charge. Profit can still fall by $\psi_{n_0}(0) - \psi_{n_0}(\tau)$.

Proof. See Appendix A.11.

Here the protected object is the marginal value h of physical stock. The auction premium can fall without changing the quantity that equates this return to marginal implementation cost. The branch comparison excludes switches to other offered capacities and exit. Its attainment premise matters: an all-offered package band has an open upper boundary, so a supremum at that boundary cannot be called its attained optimum. The full-support example in Section B.6.5 illustrates this mechanism with unchanged output 1.75 at every charge.

The maximum effect on profit and physical extraction. Let

$$\Pi^*(\tau) = \max_q V(q, \tau), \qquad \underline{\Pi} = \max_q \underline{V}(q), \tag{42}$$

where both maxima include exit and are assumed attained. The strongest selective policy leaves the gate its best surviving-market operation, which need not be the original output or the original offered capacity.

Proposition A.22 (Maximum profit loss and a quantity bound)**.** *For every charge,*

$$0 \leq \Pi^*(0) - \Pi^*(\tau) \leq \Pi^*(0) - \underline{\Pi} \leq D_0. \tag{43}$$

If the uniform cutoff (31) *is finite, then* $\Pi^*(\tau) = \underline{\Pi}$ *for all* $\tau \geq \overline{T}$*, so* $\Pi^*(0) - \underline{\Pi}$ *is the exact maximum optimized-profit loss. At every policy, any selected optimal quantity* q_τ *satisfies*

$$\Delta_0(q_\tau) \leq B_0(\tau) \leq D_0. \tag{44}$$

If, on every feasible quantity including exit, baseline profit has the quadratic separation $\Delta_0(q) \geq \kappa(q - q_0)^2/2$ *for some* $\kappa > 0$*, then*

$$|q_\tau - q_0| \leq \sqrt{\frac{2B_0(\tau)}{\kappa}} \leq \sqrt{\frac{2D_0}{\kappa}}. \tag{45}$$

Proof. See Appendix P.5.2.

The maximum contraction in physical output is, exactly,

$$\mathcal{D}_q = \sup_{\tau \geq 0} \sup_{q \in \arg\max_x V(x,\tau)} [q_0 - q]_+. \tag{46}$$

This definition includes every optimal selection at a tie. With a finite uniform cutoff, the outer supremum can be restricted to $0 \leq \tau \leq \overline{T}$. Equation (44) restricts the admissible quantities, and under quadratic separation $\mathcal{D}_q \leq \min\{q_0, \sqrt{2D_0/\kappa}\}$. Without an additional monotonicity condition on the profit ranking across quantities, the minimum output along the policy path need not occur at the saturated policy. The limiting output therefore cannot in general be used as the maximum contraction. Receipt and profit ceilings do not by themselves supply a percentage ceiling on physical extraction.

A resource bound on the largest possible production contraction. Here W is the VCG value maximized over the permitted offered capacities. The proof uses an attainable floor and uniform resource bounds; it does not require monotonicity of receipts or optimal output.

Proposition A.23 (A resource bound on policy effects)**.** *Normalize exit profit to zero. Let* $\mathcal{Q}$ *be the fixed set of attainable positive outputs and let* $\mathcal{T}$ *be a policy interval containing zero. Suppose the policy profit maximum is finite and attained for every* $\tau \in \mathcal{T}$*. Assume uniform constants* $\bar{v} > 0$ *and* $d > 0$*, and a policy-independent attainable continuation floor* F*, such that*

$$F(q) \leq W(q,\tau) \leq \bar{v}q, \qquad \mathcal{C}(q) \geq d \quad (q \in \mathcal{Q},\ \tau \in \mathcal{T}). \tag{47}$$

Define the protected profit

$$\Pi_F = \max\left\{0, \sup_{q \in \mathcal{Q}}[F(q) - \mathcal{C}(q)]\right\} > 0. \tag{48}$$

Every policy optimum is active and satisfies

$$q_\tau \geq \underline{q}_F := \frac{d + \Pi_F}{\bar{v}} > 0. \tag{49}$$

Consequently, for any baseline optimum q_0 and any $\tau \in \mathcal{T}$,

$$\frac{(q_0 - q_\tau)_+}{q_0} \leq 1 - \frac{d + \Pi_F}{\bar{v} q_0}. \tag{50}$$

The supremum defining Π_F need not be attained.

Proof. See Appendix P.5.3.

The bound is a ceiling on the possible contraction, not a claim that this contraction is attained or that optimal output falls monotonically with the policy. The full-retention option gives $F(q) = H(q)$ in the VCG auction; any stronger floor must be attainable under the auction's actual participation and trading rules. A strictly positive Π_F rules out exit. The weaker condition $\Pi_F = 0$ permits exit and therefore supplies no positive lower bound on the selected output.

When coverage can change, a common d must hold across all compared active branches. Selva's resource account supplies $\mathcal{C}_A(q) \geq d_A$; finitely many branches with $\min_A d_A > 0$ permit $d = \min_A d_A$. Without a uniform positive deficit, the bound with d_A applies only to an optimum on that branch. The nonloss inequality $q \geq d_A/\bar{v}$ is an implication of profitable operation, not an independently binding technological constraint. Neither it nor a positive acquisition margin establishes unchanged production. Exact insulation still requires the global profit comparisons, including other coverage branches and exit.

The worked example in Section B.5.1 follows the proofs.

A.3 Proof of Proposition A.2: Household choice and competitive production

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 (1); 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. □

A.4 Proof of Proposition A.3: Minimum coercion

Proof. The feasible efforts form an upper interval whose smallest element is (2). Increasing effort above it leaves the induced household allocation unchanged and increases its resource cost. The least element is therefore optimal. At points where $\Delta_v'(e_v^*) \neq 0$, implicit differentiation of the binding constraint gives the derivative. At a cap or a slack constraint, the relevant boundary condition replaces that derivative. □

A.5 Proof of Proposition 2.1: Incentives within the committed mechanism

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. □

A.6 Proof of Lemma 2.5: Whole-buyer critical payments

Proof. Telescoping the seller orders gives $H(a+j\bar{q})-H(a)$, so maximizing (17) 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. □

A.7 Proof of Proposition 2.6: Reallocation toward unaffected buyers

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. □

A.8 Proof of Proposition 2.8: A zero direct laundering response

Proof. The two equalities are the partial and total chain rules applied to (26). 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 A.2 gives an exact continuous-distribution example. □

A.9 Proof of Theorem 2.10: The exact exposure identity

Proof. Summing the pivot transfers and adding actual retention yields

$$R_n = \sum_i S_{n,-i} - (N-1)S_n \tag{51}$$

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 (51) gives the first equality in (31). 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 (32). □

A.10 Proof of Theorem A.20: Floor conditions for unchanged physical production

Proof. At fixed output, the maximum of weakly decreasing capacity values is weakly decreasing, and it is no smaller than their maximum floor. The baseline capacity n_0 remains feasible, so

$$W(q_0,\tau) \geq W_{n_0}(q_0,\tau) \geq W(q_0,0) - n_0\bar{q}\tau,$$
$$W(q_0,\tau) \geq \underline{W}(q_0).$$

Combining these inequalities gives $V(q_0,\tau) \geq V(q_0,0) - B_0(\tau)$. Every active competitor has $V(q,\tau) \leq V(q,0)$, and the same holds with equality for exit. Hence

$$V(q_0,\tau) - V(q,\tau) \geq \Delta_0(q) - B_0(\tau).$$

This proves the gap certificate and its strict version. Since $B_0(\tau) \leq D_0$, condition (37) is the corresponding all-policy certificate.

If $D_0 = 0$, the original output loses no value at any policy. Baseline optimality and weakly decreasing competitor values then preserve every baseline comparison, including its strictness. If a particular baseline capacity attains its own floor, its protected value

equals $W(q_0, 0)$; therefore the optimized floor also equals this value. Only other baseline maximizers could tie it later, proving the final assertion. The finite-gap interval follows directly from $B_0(\tau) = \min\{n_0\bar{q}\tau, D_0\}$. □

A.11 Proof of Corollary A.21: A protected marginal return and falling rents

Proof. With $H(y) = hy$, welfare can be written

$$S_n(q, \omega, \tau) = hq + \max_{\sum_i k_i \le n} \sum_i \{v_i(k_i; \omega, \tau) - h\bar{q}k_i\}.$$

The maximizing buyer allocation and each pivot payment depend on n, ω, τ but not on q within the allowed capacity branch. Subtracting the value $h\bar{q}$ of each sold package from its payment therefore yields a stock-independent premium. The opportunity floor makes their sum nonnegative, proving (39). The exact exposure identity makes ψ_n weakly decreasing, and the surviving-market limit gives $\underline{\psi}_n$. Maximizing stock within branch n gives exactly $J_n + \psi_n(\tau)$ and the same argmax quantities as at zero policy. Strict convexity of cost gives uniqueness on an interval, and the interior first-order condition is $\mathcal{C}'(q_n) = h$. The protected branch value is at least $J_{n_0} + \underline{\psi}_{n_0}$, whereas every competing branch is at most $J_n + \psi_n(0)$. Comparing with exit proves the corollary and the stated exact loss on the unchanged branch. □

A.12 Proof of Corollary A.19: The retained-value share

Proof. At fixed stock and offered capacity, the seller orders $\rho_1(q - n\bar{q}), \ldots, \rho_n(q - n\bar{q})$ and all unaffected buyer orders are fixed. Each regulated order weakly falls with the charge. In the highest-n selection with fixed tie priority, the rank of a fixed seller order can therefore only improve relative to regulated orders; its rank relative to other fixed orders does not change. Every seller order already selected remains selected. The number of retained packages, and hence y_n, weakly increases. Nondecreasing H gives a weakly increasing numerator in (32). Theorem 2.10 gives a weakly decreasing denominator. For two charges $\tau_1 \le \tau_2$ with positive receipts, write these numerators as $A_1 \le A_2$ and denominators as $B_1 \ge B_2 > 0$. Since $A_1 \ge 0$, $A_1/B_1 \le A_1/B_2 \le A_2/B_2$. When H consists only of net laundering proceeds, the same ratio is precisely their receipt share. □

B Worked examples and numerical material

B.1 Worked payment comparison

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 1 shows how unaffected competition and retention successively buffer seller value.

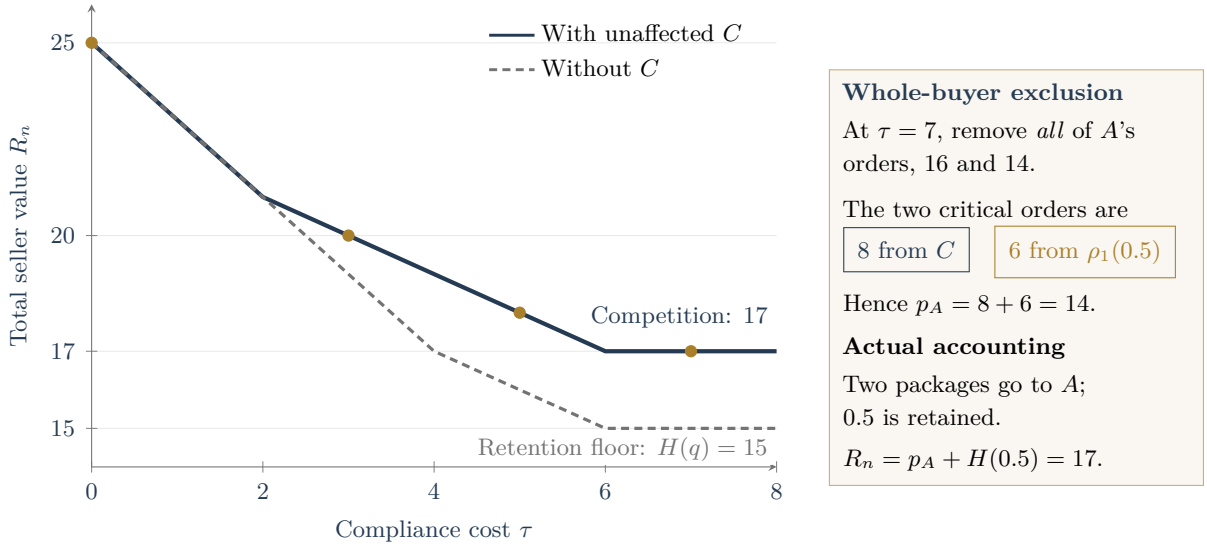

Figure 1: Two buffers for VCG seller value at fixed inventory. Dots mark the with-C observations in Table 7.

Table 7: 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

Before the policy, at $\tau = 0$, the seller's gross receipts are 25: a payment of 22 and actual retained value of 3. Buyer A's values, 16 and 14, give it both offered packages. The gate keeps the remaining 0.5 units, worth $6 \times 0.5 = 3$. To compute the VCG payment, remove *all* of A's orders and ask what the remaining buyers and the gate could obtain from the same stock. Buyer B would then receive both packages, valued at 12 and 10, while the gate would still retain 0.5 units worth 3. Their combined operating and retention value is therefore $12 + 10 + 3 = 25$. In the actual allocation with A present, and before its payment, those same participants have combined value 3: neither B nor C gets a package, and the

gate retains the residual. The VCG rule charges A the difference between these two values. Adding actual retention to that payment gives the seller's receipts:

$$p_A = \underbrace{(12+10+3)}_{\text{others' value without } A} - \underbrace{3}_{\text{others' value with } A} = 22,$$

$$R_n = \underbrace{22}_{\text{buyer payment}} + \underbrace{3}_{\text{actual retained value}} = 25.$$

The exclusion comparison uses the other participants' operating values and retention opportunities; it does not calculate the auction payments the gate would collect without A.

The intermediate rows follow the same rule. At $\tau = 3$, B's package values fall to 9 and 7. Without A, one package would go to B at value 9 and one to C at value 8; subtracting the unchanged residual value from $9+8+3$ gives a payment of 17, and receipts are $17+3 = 20$. At $\tau = 5$, the two replacement values are 8 from C and 7 from B, giving a payment of 15 and receipts of $15 + 3 = 18$.

At $\tau = 7$, B's values have fallen to 5 and 3. Buyer A still receives both packages, but its exclusion comparison now changes: one package would go to C, worth 8, and the gate would retain the other package together with the residual. Those 1.5 units have retention value $6 \times 1.5 = 9$. The remaining participants could therefore obtain combined value $8+9 = 17$ without A, compared with 3 when A wins. Its payment is their displaced value:

$$p_A = (8+9) - 3 = 14 = \underbrace{8}_{\text{unaffected buyer}} + \underbrace{6}_{\text{one extra retained package}}.$$

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.

After the policy, at $\tau = 7$, the actual allocation still gives both packages to A and leaves only 0.5 units with the gate. The seller receives 14 in cash and retains value 3, so its gross receipts are $R_n = 14 + 3 = 17$, down from 25 before the policy.

B.2 Procurement and unchanged-production example

Different buyers, the same production. The procurement example in Appendix A.2, with profits given by (55), makes the distinction concrete. It has four buyers, each able to take one unit package. One buyer's value falls with the compliance charge from an initial value of 10; the unaffected values are 9, 8 and 7. Retention has zero value, all stock is offered, and the stipulated implementation cost is 2.2 times the square of the package count. This is a downstream implementation-cost illustration, not a calibration of Selva's

household primitives. The appendix compares all feasible quantities, from zero to four packages.

Compliance cost	Buyers receiving gold	Packages produced	Seller receipts	Profit
0	Regulated and highest unaffected	2	16	7.2
2.5	Two highest unaffected	2	15	6.2

Between these two observations, the regulated buyer is replaced by an unaffected buyer, receipts fall, and profit falls. Yet two packages remain the unique optimal production choice. The buyer switch occurs as the charge passes 2; production switches to one package only as it passes 2.7, with the relevant choices tied at each threshold. Reallocation can therefore occur without a production shift, even without complete protection of the seller's rent.

B.3 A positive spread need not cover procurement

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.4 Liquidity and the retention reserve

B.4.1 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.5 Who buys the gold and who sets its payment

B.5.1 Who buys the gold and who determines its payment

Table 8 separates who buys the gold from who determines its payment. Five distinct buyers can each purchase one package. Unaffected buyers U_1, U_2, U_3, U_4 value it at $12, 11, 10, 8$; regulated buyer R bids $9 - \tau$. Feasible production is $q \in \{0, 1, 2, 3, 4\}$, all produced packages are offered, and full implementation cost is $\mathcal{C}(q) = 0.8q^2$. The gate can profitably launder one package, with net return 13; additional retained stock has no value, so $H(y) = 13 \min\{y, 1\}$.

Panel A holds the original stock at four packages. The same three unaffected buyers purchase throughout, yet R initially determines what they pay without buying any gold. Once U_4's bid replaces R's as the competing offer, receipts reach their floor; removing R then has no further effect at this stock. Panel B allows production to adjust. Selling fewer packages makes U_3 the competing buyer and raises the payment per purchaser even as total receipts fall. The gate fully uses its profitable laundering capacity in every case; laundering therefore accounts for a larger share of the remaining receipts.

Table 8: Who buys the gold and who determines its payment

	No policy $\tau = 0$	Partial policy $\tau = 2$	R removed $\tau \to \infty$
A. Payments at the original stock, $q = 4$			
Purchasing buyers	U_1, U_2, U_3	U_1, U_2, U_3	U_1, U_2, U_3
Packages received by R	0	0	Excluded
Buyer determining payments	R	U_4	U_4
Payment per purchasing buyer	9	8	8
Total buyer payments	27	24	24
Laundered packages	1	1	1
Net laundering return	13	13	13
Seller receipts	**40**	**37**	**37**
B. Production chosen by the gate			
Production	**4**	**3**	**3**
Purchasing buyers	U_1, U_2, U_3	U_1, U_2	U_1, U_2
Buyer determining payments	R	U_3	U_3
Payment per purchasing buyer	9	10	10
Laundered packages	1	1	1
Seller receipts	40	33	33
Implementation cost	12.8	7.2	7.2
Profit	**27.2**	**25.8**	**25.8**
Share of original profit	100%	94.9%	94.9%
Maximum receipt loss at original stock $q = 4$			$40 - 37 = 3$
Maximum optimized-profit loss			$27.2 - 25.8 = 1.4$
Maximum production contraction			$(4 - 3)/4 = 25\%$
Minimum production across tariffs			3 packages

Notes. Each purchasing buyer receives one package. Regulated buyer R receives none in either panel. Panel A fixes stock; Panel B compares every feasible quantity, including exit. With $d = 0.8$, $\underline{\Pi} = 25.8$ and $\bar{v} = 13$, the production floor is $q_\tau \geq 133/65 > 2$, hence at least three indivisible packages. Production switches from four to three at $\tau = 7/15$, with both quantities optimal at the threshold. The profit floor preserves $25.8/27.2 \simeq 94.9\%$ of original profit.

B.6 Analytical transmission and procurement examples

B.6.1 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

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

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.6.2 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

$$T_{\text{cell}} = \inf_{\substack{q \in I,\ a,\ k \text{ feasible in } a \\ m_{\mathcal{R}}(k_a^0) > m_{\mathcal{R}}(k)}} \frac{F_a(k_a^0, q, 0) - F_a(k, q, 0)}{\bar{q}[m_{\mathcal{R}}(k_a^0) - m_{\mathcal{R}}(k)]} \tag{52}$$

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 2.13. 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

$$R_1(q, \omega, \tau) = hq + [\min\{U, R - \bar{q}\tau\} - h\bar{q}]_+ \tag{53}$$

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

$$W_1(q, \tau) = hq + \psi(\tau), \quad \psi(\tau) = \int_{h\bar{q}}^{\infty} [1 - F_U(t)][1 - F_R(t + \bar{q}\tau)]dt \tag{54}$$

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.6.3 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

$$\Pi_1(\tau) = \begin{cases} 6.8, & 0 \le \tau \le 1, \\ 7.8 - \tau, & 1 \le \tau \le 2, \\ 5.8, & \tau \ge 2, \end{cases}$$
$$\Pi_2(\tau) = \begin{cases} 7.2, & 0 \le \tau \le 2, \\ 11.2 - 2\tau, & 2 \le \tau \le 3, \\ 5.2, & \tau \ge 3. \end{cases} \tag{55}$$

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.6.4 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 $1 \le q < 2$, the seller's reserve is $\rho(q) = H(q) - H(q-1) = 12 - 4q \le 8 < 9$, so the unaffected buyer wins. Its payment is $\max\{8 - \tau, \rho(q)\}$. Adding retained value gives

$$W(q, \tau) = \max\{H(q), 8 - \tau + H(q-1)\} \tag{56}$$

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:

$$H(q) - 2q^2 = 6.25 - 4(q - 1.25)^2,$$
$$8 - \tau + H(q-1) - 2q^2 = 8.25 - \tau - 4(q - 1.75)^2$$

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.6.5 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 (54) gives

$$\psi(\tau) = \begin{cases} (3-\tau)^2(6+\tau)/600, & 0 \leq \tau < 3, \\ 0, & \tau \geq 3. \end{cases} \tag{57}$$

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.6.6 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

$$W_1(q, \tau) = H(q) + \int_{\rho(q)}^{10-\tau} (1 - t/10)(1 - (t + \tau)/10)dt \tag{58}$$

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

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

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

$$\gamma = .04u_0 = 1 - \frac{\sqrt{19}}{5} \in (0, 1), \qquad \Lambda = (8 - .32u_0)^{-1} > 0$$

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.6.7 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

$$\Pi(q, \tau) = \begin{cases} 26q - 8q^2 - 12, & 0 \leq q \leq a_\tau, \\ -4q^2 + (14 - 2\tau)q + (6 + \tau)^2/4 - 12, & a_\tau \leq q \leq 2. \end{cases} \tag{59}$$

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

$$\bar{\chi}(q,\tau) = (6 - \tau - \rho(q))/2, \qquad \gamma = \partial_q \bar{\chi} = 2, \qquad \Pi_{qq} = -8$$

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.

Mining and governance: selected observations in 2023

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





South America

A Colombia: authorization

Unauthorized share of mapped on-land alluvial-gold footprint, 2023.

75.8%

0 100%

79,631 / 105,060 hectares
National monitored footprint.
Area, rather than mineral output.

B Brazil: mining footprint

Garimpos' share of mapped mining area, all minerals, 2023.

64.4%

0 100%

283,791 / 440,416 hectares
National area; MapBiomas Collection 9.
Classification by activity type.

Africa

C Ghana: protected forests

Forest reserves affected by illegal mining, as of May 2023.

11.8%

0 100%

34 / 288 forest reserves
National administrative count.
Affected reserves, rather than the share of forest area destroyed.

D Eastern DRC: armed interference

Estimated workers at sites with armed interference / workers at all visited sites.
Visits: 21 Jan.–20 June 2023.

Gold **87.1%**
24,014 / 27,579 workers; 190 sites

3T **53.4%**
1,896 / 3,553 workers; 36 sites

0 100%

Purposive site sample; 35 of 36 active 3T sites are in Sud-Kivu. Mixed sites excluded.

Figure 2: Selected 2023 observations describe distinct dimensions of mining and governance; the panels do not estimate a common prevalence or trend.

Sources: A, UNODC/MME, pp. 22, 48; B, MapBiomas Collection 9, March 2025 workbook, 2023 column; C, Ghana 2025 VNR, p. 128, May 2023 observation; D, authors' calculations from IPIS on-site records dated 2023.

3T: tantalum, tin and tungsten minerals. DRC coverage differs across minerals and is not nationally representative; the shares do not identify a mineral effect. No sampling-based confidence intervals are implied.

Map boundaries: Natural Earth; equal markers locate countries only. Exact values, definitions, data versions and exclusions are documented in the companion empirical source note. The separate audit `data/source_audits/D9_observed_evidence.md` records related evidence checks and their limits.

Part II

Section 4: certification architecture

C Supporting proofs and technical definitions

C.1 Certification: general reference account and worked example

This reference account retains the physical, contractual and incentive definitions needed to read the proofs. Its general admission discussion includes admission away from extraction. The paper's intended mine-gateway application is described in the opening note; the broader account is retained here as a reference variant. The full posted-contract example ends with live contamination 3/23, approximately 13.04%. It is an exact worked example, distinct from the saved gold comparison reported in the numerical appendix. Clean procurement, already acquired dirty stock, outside deliveries, exclusion and resource costs all remain in the example's account.

C.1.1 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

$$v(O) + m_{\text{loss}} = v(I), \qquad \psi(O) \subseteq \psi(I), \qquad m_{\text{loss}} \geq 0, \tag{60}$$

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

$$0 \le d_k \le q_{D,k}, \qquad k \notin \mathcal{I} \implies q_{D,k} = d_k = 0. \tag{61}$$

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.

C.1.2 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

$$C_t = \{c \in S_t : \text{Anc}(c) \cap (\mathcal{O}_t^{\text{rev}} \cup \mathcal{O}_t^{\text{susp}}) = \varnothing\}. \tag{62}$$

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 C.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*

$$v(S_t) = m\left(\bigcup_{c\in S_t} \psi(c)\right) = v(S_0) + A_t - X_t, \tag{63}$$

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. See Appendix P.6.1.

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:

$$Q_R(t) = \sum_{c\in S_t} v(c)\mathbf{1}\{\mathrm{Anc}(c) \cap \mathcal{O}_t^{\mathrm{rev}} \neq \varnothing\}. \tag{64}$$

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 (62) 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,

$$\begin{aligned} \iota_{\text{world}}^{\text{live}} &= \frac{\sum_k \mathbb{E}[U_k]}{\sum_k \mathbb{E}[V_k]} = \sum_{k \in \mathcal{I}} w_k \iota_k^{\text{live}} \leq \sum_{k \in \mathcal{I}} w_k, \\ w_k &= \frac{\mathbb{E}[V_k]}{\sum_h \mathbb{E}[V_h]}, \qquad \iota_k^{\text{live}} = \frac{\mathbb{E}[U_k]}{\mathbb{E}[V_k]}. \end{aligned} \tag{65}$$

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.

C.1.3 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.

C.1.4 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

$$A_j = P_{-j}L_{L,j} + (1 - P_{-j})L_{U,j} - L_{E,j}. \tag{66}$$

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 C.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*

$$\pi_j s_j(e_j)A_j - c_j(e_j). \tag{67}$$

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

$$c_j'(e_j) = \pi_j A_j s_j'(e_j). \tag{68}$$

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

$$\frac{de_j}{dA_j} = \frac{\pi_j s_j'(e_j)}{c_j''(e_j) - \pi_j A_j s_j''(e_j)} > 0 \tag{69}$$

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 (67) covers f_j.

Proof. See Appendix P.6.2.

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, (67) 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.

C.1.5 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

$$P_n = 1 - (1 - r_0) \prod_{j=1}^{n} (1 - r_j). \tag{70}$$

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 C.3 (Conditional survival of an admitted fraud)**.** *Under the conditional independence and effective-decision requirements, adding a channel changes discovery probability 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. See Appendix C.2.

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 (70) becomes

$$P_n = 1 - (1 - r_0) \exp\left(-\alpha \sum_j e_j\right). \tag{71}$$

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 (70), 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.

C.1.6 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 $0 < d(a) \le q$. 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

$$G_\Omega(a) = F_C(a;\Omega) - k_f(a;\Omega) - \Big[W_D(q;\Omega) - W_D(q - d(a);\Omega)\Big]. \tag{72}$$

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.

C.1.7 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 (72) 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

$$\mathcal{B}(T) = \sum_{i \in S \setminus \{0\}} \frac{a_i(T)}{\lambda_i} + \sum_{j \in T} \frac{\tau_j(T)}{\lambda_j}, \tag{73}$$

$$K_J^F(a) = \min_{\substack{T \subseteq J \\ \mathcal{B}(T) \leq F_0(a,\Omega)}} \{a_0(T) + \mathcal{B}(T)\}, \tag{74}$$

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 C.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 $G_\Omega(a) \le K_J^F(a)$ for every feasible nonzero joint action. Strict inequalities $G_\Omega(a) < K_J^F(a)$ exclude indifferent implementations as well.*

Proof. See Appendix P.6.3.

Proposition C.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*

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

Then $K^F_{J\cup\{j_{\text{new}}\}}(a) \ge K^F_J(a)$.

Proof. See Appendix C.3.

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.

C.1.8 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 $D_i \ge 0$. Put

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

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

$$K_n = \min_{T\subseteq\{1,\dots,n\}} \left\{ \sum_{j\in T} c_j^{\text{sil}} + D_G P(T) \right\}. \tag{75}$$

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 $P(T)\sum_{i\neq 0} D_i/\lambda_i +$

$\sum_{j\in T} c_j^{\mathrm{sil}}$ 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 C.6 (Homogeneous channels and saturation)**.** *If $r_j = r \in (0,1)$ and all channels have the same gross silence cost $c_j^{\mathrm{sil}} = c^{\mathrm{sil}} \geq 0$, then*

$$K_n = \min\{D_G P_n,\ nc^{\mathrm{sil}} + D_G r_0\}, \qquad P_n = 1 - (1-r_0)(1-r)^n. \tag{76}$$

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

Proof. See Appendix C.4.

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 (74) 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

$$u \leq u_{\max} := \left\lfloor \frac{\log[(1-\delta)/(1-r_0)]}{\log(1-r)} \right\rfloor. \tag{77}$$

If each suppressed channel costs at least $c_{\min}^{\mathrm{sil}} > 0$, maintaining this target costs at least $c_{\min}^{\mathrm{sil}} \max\{0, n - u_{\max}\}$. 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.

C.1.9 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

$$\iota^{\text{in}} = \frac{\mathbb{E}U}{\mathbb{E}A}, \qquad \iota^{\text{live}} = \frac{\mathbb{E}(U - R_U)}{\mathbb{E}(A - R_U - R_A)}, \tag{78}$$

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 C.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*

$$\iota^{\text{live}} = \frac{\iota(1 - P)}{1 - \iota P} \leq \iota, \qquad \iota - \iota^{\text{live}} = \frac{\iota P(1 - \iota)}{1 - \iota P}. \tag{79}$$

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:

$$\frac{\partial \iota^{\text{live}}}{\partial \iota} = \frac{1 - P}{(1 - \iota P)^2}, \qquad \frac{\partial \iota^{\text{live}}}{\partial P} = -\frac{\iota(1 - \iota)}{(1 - \iota P)^2}. \tag{80}$$

Proof. See Appendix C.5.

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.

The worked example is in Online Appendix D.3.

In an ongoing market, (78) 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.

C.1.10 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, 2021a]. 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

$$\begin{aligned}\phi V + E_0 \;\geq\; & C_{\text{admission}} + C_{\text{custody}} + C_{\text{records}} + C_{\text{forensic}} + C_{\text{adjudication}} \\ & + C_{\text{availability}} + C_{\text{capital}} + C_{\text{promised obligations}}.\end{aligned} \tag{81}$$

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 C.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*

$$V \leq \frac{\varphi I + \bar{E}_0 - F}{c} \tag{82}$$

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

Proof. See Appendix C.6.

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.

C.1.11 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.

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

$$G(c) = \Gamma - c, \qquad \Gamma = h - p_D. \tag{83}$$

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.

Joint solution and disclosure incentives Define $A = rL_L - L_E$ and impose the strict condition

$$A > (1-r)D. \tag{84}$$

Condition (84) makes truthful disclosure privately preferable to any silence agreement the original coalition can finance. The bargaining argument is given in Appendix P.6.4.

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

$$R(\pi, e) = \frac{\kappa e^2}{2} + \pi\{eL_E + (1-e)rL_L\} = \pi r L_L - \pi A e + \frac{\kappa e^2}{2}. \tag{85}$$

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

$$p(e) = e + (1-e)r = r + (1-r)e. \tag{86}$$

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 C.9 (Positive-trade interior equilibrium)**.** *Maintain* (84)*, let* $a = A/\kappa$*,* $T = \Gamma - rD$*, and assume*

$$0 < T < C, \qquad a > 0, \qquad aT < C + D(1-r)a. \tag{87}$$

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

$$\pi^* = \frac{T}{C + D(1-r)a}, \qquad e^* = a\pi^*, \qquad p^* = r + (1-r)e^*. \tag{88}$$

The unit mass of slots trades; $0 < \pi^* < 1$*,* $0 < e^* < 1$*, 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. See Appendix P.6.4.

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

$$v > h + \phi + R(\pi^*, e^*). \tag{89}$$

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.

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

$$\iota(\pi, p) = \frac{\pi(1-p)}{1-\pi p}. \tag{90}$$

Continuing scrutiny lowers live contamination: $\iota(\pi^*, p^*) < \iota(\pi_0, r)$, accounting for the smaller surviving denominator as well as fewer illegal admissions. The derivative comparison is in Appendix P.6.4.

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

$$\int_0^{C\pi^*} \frac{\Gamma - p^*D - c}{C}\, dc = \frac{C(\pi^*)^2}{2}. \tag{91}$$

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

$$\mathcal{C}_0 = \phi_0 + \pi_0 r L_L, \qquad \mathcal{C}_1 = \phi_0 + H + R(\pi^*, e^*). \tag{92}$$

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.

The worked example is in Online Appendix D.4.

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 C.1.1 determine aggregate leakage. The country composition of actual supply consequently matters for both remaining exposure and the incidence of prevention costs.

C.2 Proof of Proposition C.3: Conditional survival of an admitted fraud

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. □

C.3 Proof of Proposition C.5: A funded monotonicity condition

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. □

C.4 Proof of Proposition C.6: Homogeneous channels and saturation

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

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

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 $c^{\mathrm{sil}} > 0$ the second eventually exceeds D_G, whereas for $c^{\mathrm{sil}} = 0$ it is the smaller endpoint. The case $n=0$ follows directly. □

C.5 Proof of Proposition C.7: The removal effect

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 is sufficient. □

C.6 Proof of Proposition C.8: Cost of prevention and a scale limit

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. □

D Worked examples and numerical material

D.1 Mixed-lot withdrawal example

Whole-lot removal can behave differently: from 100 units with 10 illegal, withdraw a mixed lot with 9 illegal and 89 admissible units. The surviving stock has 1 illegal unit among 2. Illegal mass falls, but its share rises from 10% to 50%. The Online Appendix's *Admission, retrospective removal, and surviving leakage* gives the general result and boundary cases.

D.2 Saved gold calibration

Saved gold calibration. Table 9 puts surviving contamination at approximately 3.5 false-origin units per 10,000 live certified units, falling to approximately 1.5 under stronger late discovery. Buyer effort and fraud adjust in the saved equilibrium, while the admission environment and expenditure backing remain fixed. This differs from the preceding comparison that removes early scrutiny.

Table 9: False-origin material among live certified gold: a saved calibration.

	Reference	Stronger discovery
Independent late-discovery parameter	0.25	0.50
False-origin units per 10,000 live certified units	3.47	1.47

Notes: Gold, VCG, central saved scenario `eta1_unknown0.5_retail0.25`: bulk supply elasticity 1, unknown-origin candidate share 0.5, retail share 0.25, without withholding; 10% vulnerability at named eligible origins and zero in the rest of the world. The rate divides surviving false-origin mass by live certified mass after withdrawals. Buyer effort and fraud respond within the saved equilibrium. These are calibrated model levels, not an observed fall in illegal extraction. *Source:* Frozen C1 baseline and screening states, exported in `paper-reference-data.json`, central gold–VCG comparison.

The paper retains 15–20% as a working benchmark for illegal gold in current supply. That denominator differs from live certified supply; the benchmark is not a measured initial condition for this calibration. The saved comparison does not quantify the later production transition or incorporate a launch-stock or recycled-supply extension. Full inputs, physical masses and sensitivity, together with the separate exact 13.04% illustration, belong to the Online Appendix.

D.3 Whole-lot withdrawal and the surviving denominator

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 C.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.

D.4 Lower leakage and higher funded cost

Example D.1 (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 $F = 1$ is sufficient. 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.

Part III

Sections 5–6: markets, research and government

E Current asymmetric growth and mixed-site policy: proofs

These proofs support the revised Sections 5–6. The earlier one-period certified-innovation, fixed-expenditure and owned-installation models are preserved with their exact sources in the frozen pre-revision appendix; they are comparison specifications, not proofs of the baseline here. Machine and research ownership, physical implementation and dated finance must satisfy the hypotheses stated below. Conditional analytical results are distinguished from the still-required full-market and Lean adapters in the accompanying result and evidence map.

E.1 Result preservation, model comparisons and proof status

The maintained baseline is persistent certified patents and one-period dirty appropriation followed by diffusion, within a fixed scientist pool. The following results are analytical conditional results. A condition in this table is part of the theorem, not a fact established by its label or by an older Lean compilation.

Required result	What is retained or strengthened	Exact boundary and evidence
Certification price shock	Main mixed-site theorem gives the unique price root and a sharp demand threshold; the inherited closed form remains a benchmark.	Certification-date budgets and technology are held fixed for the comparison. Common-quote receipt equivalence is required. Gross certified price and net receipt have separate tests. Analytical proof: Section E.3.
Profitability divergence	Certified operating and research opportunities grow proportionally to certified technology on the stated path; actual dirty opportunities have a sublinear bound.	Endogenous wage and income, positive outside dirty capacity, compatible regional supply, licensed core, finite network bounds and actual claim envelopes. Absolute dirty receipts can grow. Patent, company and fiscal claims remain separate.

Required result	What is retained or strengthened	Exact boundary and evidence
Fixed science	The inherited success functions, convex laboratory costs, owner masses and marginal allocation are retained.	Verification must actually fit within S; graduation only changes eligibility. The older linear payoff-level scientist-simplex theorem is not the nonlinear owner-effort proof.
Support threshold	The minimum additive funded grant is retained with new, correctly owned patent and dirty values.	The exact marginal-return formula uses the raw certified prize. Zero rewards and zero costs permit ties. General-equilibrium policy changes require reward revaluation. The new registered Lean module verifies the threshold algebra only.
Path dependence	An explicit sufficient cutoff, a robust private-green region and temporary-support hitting conditions remain main-text results.	Actual unsubsidized equilibrium must remain in the stated domain. Deterministic hitting requires the exact cohort law; stochastic hitting and its finance are distinct. Continuing support and failed-withdrawal cases are included.
Optimal taxation	The original profitable-opening ceiling is preserved; shared HQ changes the gross fiscal elasticity.	Exact net fiscal value, parent participation, tax commitment and dated cash. Fixed-HQ closed form and responsive-HQ stationary condition have distinct scopes.
Intervention	The optimized feasible control value enters a Bellman problem with waiting, failed attempts and next-period success.	Normalized contraction and strict fiscal margin; actual finance on every compared history. Positive future NPV does not supply current reserves.
Conservation game	Governments interact through endogenous markets and joint company networks; a regular conditional equilibrium theorem is supplied.	The full auction/stage response, complete funded feasible-plan continuity and endogenous network variants are not established by the fixed-point theorem itself.
Physical conservation	Individual success, mass decline, critical-site safety and cumulative damage each have a separate claim.	Actual legal safety and growing dirty intensity matter. $qg_D < 1$ bounds the disturbance tail; a finite critical set and a continuum have different all-site conclusions.

The logical connections. The diagram distinguishes economic inputs from the conditions needed to apply the dynamic results. A missing arrow is a proof obligation, not a new institution introduced to force the conclusion.

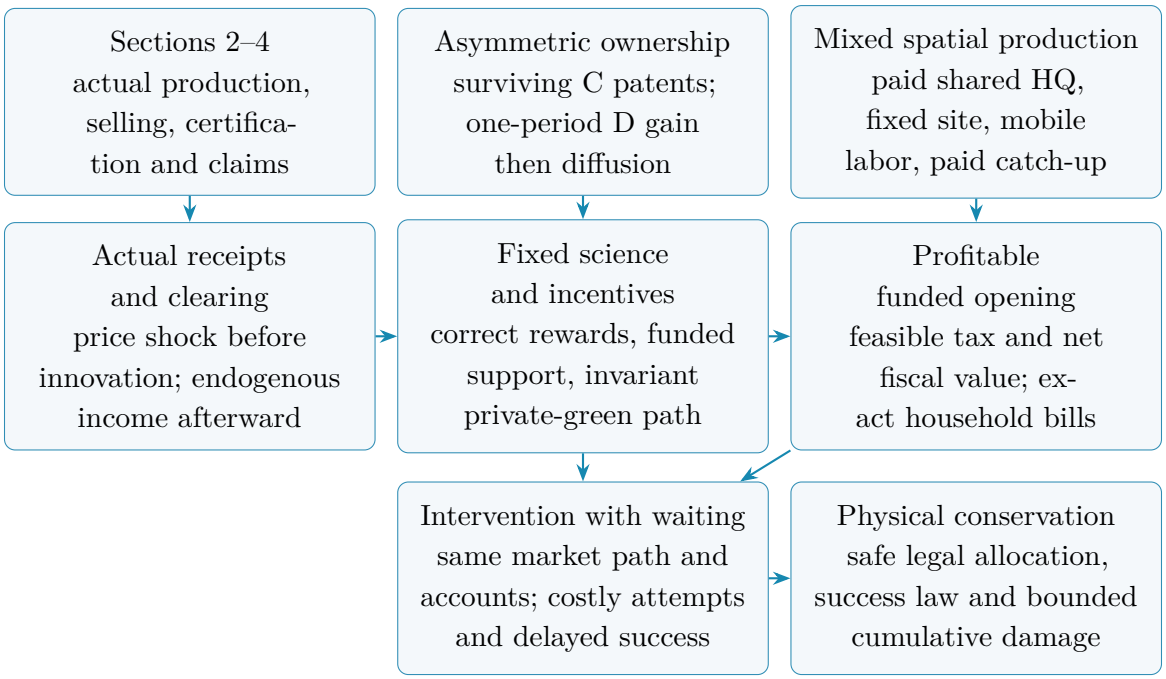

Explicit alternatives, not ambiguity in the baseline.

Comparison	Consequence for rewards or production	Status relative to this revision
One-period certified rights	Only the next productive-period gain enters, plus later contractual differences.	Historical comparison; its old affine prize cannot replace the surviving-patent baseline.
Persistent dirty patents	Dirty reward requires the same patent's survival-weighted future profits, not sector-wide future frontier profits.	Alternative comparison. In a benchmark with old-patent flow bounded by $G_A^{\theta n}$, $\delta_R G_A^{\theta} < 1$ is sufficient for summability; survival can weaken that condition. The present one-period dirty theorem does not need this annuity.
Immediate shared dirty knowledge	With no owned productive advantage or contractual gain, positive appropriation need not exist.	It cannot be given an invented monopoly prize. This differs from the maintained one-period advantage followed by sharing.
Fixed regional workforces	Labor allocation and the CES price bounds change.	A separate benchmark can give $\theta = [1 + \beta_m(\varepsilon_M - 1)]^{-1}$. It is not the mobile fixed-site exponent θ_Σ.
Mobile labor without a productive fixed site	The old relative-wage condition $w_J/w_C = \chi x(1-\tau)^{1/\beta_m} < 1$ can empty local legal production.	Obstruction retained. The mixed model derives positive hiring at an operating site from diminishing returns, while opening remains a choice.

Comparison	Consequence for rewards or production	Status relative to this revision
Endogenous corporate delay	The parent can wait after control, with public support bills accruing during delay.	Optional extension. The baseline's funded feasible opening at control is not a solved corporate stopping problem.

Aggregation that is retained and information that cannot be discarded. The homogeneous regional cost moment remains exact within its inherited installed-input specialization. The mixed economy additionally needs the company membership and the distinct moments $\mathcal{J}_j, \mathcal{K}_j$. The sector's average technology does not recover the quality distribution of a particular inventor's assigned line. The fixed scientist account does not recover that ownership either. Legal-entry values retain catch-up age, and subsistence bills retain the earnings distribution of eligible households. Dated balances and historical claim cohorts are needed when finance or verification binds. Losing these distinctions would create spurious profits or omit real public costs.

What remains to be proved for the complete economy. The linear-retention case now derives a common producer receipt directly from the inherited whole-buyer auction, under its stated value and delivery conditions. This local bridge does not yet prove the full finite-auction and funded certification correspondence has those receipts, nor that every allowed opening, research and tax deviation stays in its regular domain. The date-specific North–South price comparison and the CES production growth application are explicitly separate closures; a single fully funded North–South transition must connect their demand and income accounts before their numerical thresholds are combined. A complete endogenous network, entry and government equilibrium requires its own construction. The specified public/private funding contracts, uniform verification capacity, licensed-market bounds and physical legal-site implementation must also be verified for that application. These obligations are retained; no result is recovered by editing Sections 2–4 or by assuming permanent green research.

Lean reuse at actual hypotheses. Machine demand and feasible supplier optimization, surviving-patent sums, cohort entry recursions, Bellman limits, dated cash and success bounds are reusable registered components. The new `AsymmetricSupport` module checks the marginal support algebra, the zero-cost clipping boundary and abstract invariance/reachability implications. It was compiled through the current wrapper, immediately registered, audited and independently reviewed. Its assumptions do not establish the corresponding economic hypotheses. The current nonlinear owner-effort allocation, corporate/site optimization, actual tax and market adapters, robust economic domain, stochastic stopping and full equilibrium remain separate formalization tasks. The evidence

map reads the shared register; it does not create a second register or confer manuscript acceptance.

Economic meaning and contribution. In ordinary terms, a legal mine can hire mobile workers because the deposit itself is productive: workers do not have to be trapped there. A company with an existing network can spread paid headquarters services across its branches. That changes both what a new mine is worth to the company and how much the host can tax without preventing it from opening. Research changes those opportunities, but scientists still compare the rewards they can actually keep. A temporary subsidy can end only when the ensuing green path remains attractive on its own. Governments act when the funded legal opportunity outweighs protection income and intervention costs. Conserving the forest adds a further requirement: the resulting physical activities must be safe at the relevant sites, soon enough to respect the damage budget.

The proposed contribution links corporate scope and local fixed resources to directed innovation, feasible taxation and state-building under mobile labor. It combines ingredients with clear antecedents in Kleinman, Aghion–Howitt and spatial growth. The conditional comparative statics are proved here; a priority claim or a completed unrestricted spatial growth equilibrium is not implied.

E.2 Production, actual receipts and aggregation: current proofs

The receipt bridge. The inherited implementation correspondence gives the implementable quantities $\mathcal{Q}_r$ and minimized procurement cost $\mathcal{C}_r$. Its attained maximization is therefore exactly (94). No averaging or reallocation of buyers is involved. If selling continuation is differentiable and concave, the interior first-order condition is $W_r'(q) = \mathcal{C}_r'(q)$, with the appropriate one-sided inequalities at boundaries. For a machine-using operator, the chain rule similarly gives $W_r'(Q)\partial Q/\partial x^M = q^M$. If W_r' varies with Q, changing a supplier's quote also changes this marginal receipt. Consequently the constant-receipt elasticity used in the monopoly calculation below cannot be asserted for nonlinear W_r. Whole-buyer pivot transfers and retained stocks retain their inherited definitions; neither aggregate welfare nor an outside price replaces them.

An exact retention-dominant whole-buyer subclass. Suppose $H(y) = p_D y$ on the relevant stocks and every buyer's marginal package value is at most $p_D\bar{q}$, with zero value at zero allocation. Then every feasible auction allocation has welfare at most $p_D q$, while retaining all stock attains $p_D q$. Therefore $S_n = S_{n,-i} = p_D q$ for every buyer and every offered capacity. If an efficient allocation awards buyer i any packages, their joint value must equal $p_D\bar{q}k_i$; a strict inequality would make retention more efficient. The Clarke pivot payment is consequently $p_i = S_{n,-i} - S_n + v_i(k_i) = p_D\bar{q}k_i$, even for whole-buyer blocks. Adding the retained stock value yields $R_n = p_D q$, hence $W_r(q) = p_D q$. The proof accommodates the inherited commitment to offer every complete package and arbitrary efficient allocation of ties. It does not replace any Section 3 rule. The existing

laundering/storage choice must actually implement the stated H. In the zero-channel-cost laundering specialization with sufficient capacity and storage return no greater than $p_D y$, the outside payment $\theta_E = p_D$ implements it. To identify this retained output with current final-market supply, actual terminal buyers must fund delivery at p_D, and retained-use sales and auction sales must be counted as disjoint dispositions of the same stock. A continuation storage value alone supplies neither current cash nor current consumption. This local receipt identity therefore strengthens the common-quote application but does not establish a general full-world auction, trading or funding equivalence.

Machine demand and admissible monopoly pricing. Fix positive p, Φ, line quality a, and a linear machine quote $q^M > 0$. The linewise problem is

$$\sup_{x \geq 0} \left\{ \frac{p \Phi^{\beta_m} a^{\beta_m}}{1 - \beta_m} x^{1-\beta_m} - q^M x \right\}.$$

Its derivative decreases from infinity to $-q^M$, so strict concavity gives $x = a(p/q^M)^{m_\beta}\Phi$. Substitution and integration give (96). In particular, machine expenditure is $(1-\beta_m)pQ$, and $U^{op} = \beta_m pQ$. When the patent supplier has marginal cost $c_M = 1 - \beta_m$, it chooses

$$(q^M - c_M) a (p/q^M)^{m_\beta} \Phi.$$

Its derivative has the sign of $m_\beta c_M - (m_\beta - 1)q^M$, hence the unique positive optimum over all nonnegative quotes is $q^M = 1$. The value tends to zero as the quote tends to infinity and is negative below cost. This proof requires that these quotes, sales and ownership rights are admissible and that the supplier takes p, Φ as given; strategic effects on the final mineral price are excluded by the price-taking line-market specialization. The successful line's unconstrained margin is $\beta_m a p^{m_\beta} \Phi$.

This unconstrained calculation must also pass predecessor competition. For interchangeable vintages of one line, effective machine service is $a^{\beta_m/(1-\beta_m)}x$; its unit cost is $q^M / a^{\beta_m/(1-\beta_m)}$. Let the leading quality be $a = (1+\gamma_C)a_{old}$ and allow the displaced predecessor to be copied competitively at c_M. The leading design supplies the line strictly when

$$\frac{q^M}{a^{\beta_m/(1-\beta_m)}} < \frac{c_M}{a_{old}^{\beta_m/(1-\beta_m)}}, \qquad q^M < \bar{q}_C^M = c_M (1+\gamma_C)^{\beta_m/(1-\beta_m)}.$$

Thus the maintained application $\bar{q}_C^M > 1$, including inherited lines and excluding a stronger rival alternative, places the unconstrained maximizer strictly within the actual sales-preserving domain. Above the limit quote sales vanish; below it the objective is the one already optimized. Quote one is therefore an actual global optimum without a franchise blocking public copying. Equivalently, the quality step satisfies $\gamma_C > (1-\beta_m)^{-(1-\beta_m)/\beta_m} - 1$.

At a limit quote at or below one, an allocation rule awarding quality-cost ties to the leading design gives the constrained maximum $q^M = \min\{1, \bar{q}_C^M\}$. If ties do not award all demand to the leader, arbitrarily close undercutting can approach the value at the cap without attaining it. Existence and allocation at that boundary must then be specified. In

the capped-quote economy, all formulas use the actual quote: the operating coefficient is multiplied by $(q^M)^{-(m_\beta-1)}$, supplier profit is $(q^M-c_M)a(p/q^M)^{m_\beta}\Phi$, and the supplier-profit and manufacturing shares of net mineral sales are respectively $(1-c_M/q^M)(1-\beta_m)$ and $(c_M/q^M)(1-\beta_m)$. When caps differ across lines, both output and profit require the corresponding quality-weighted quote moments. The unit-quote company and vintage formulas cannot be retained merely by relabeling the patent margin.

Public copying instead gives $q_D^M = c_M$. The demand formula still holds, the operating coefficient is multiplied by $c_M^{-(m_\beta-1)}$, and supplier margin is exactly zero. It would be incorrect to keep the monopoly coefficient or add patent income after making that line freely copyable. A completely integrated parent that owns and jointly chooses all machines also pays real marginal costs; its internal invoices cancel, so its research reward must be its actual portfolio gain or its inventor contract. That is a distinct ownership comparison, not the independent-supplier baseline.

Local hiring and accounting. Conditional on $H>0$, write $v = a^{op}(H^{site})^{1-\alpha_\Sigma}H^{\alpha_H} > 0$. The local objective is $vL^{\alpha_L} - w^L L$. Its unique maximizer is $L = (\alpha_L v/w^L)^{k_L}$; its first-order condition gives $w^L L = \alpha_L v L^{\alpha_L} = \alpha_L U^{op}$. This proves (98). For an independent certified machine supplier, net mineral sales are $V = (1-\tau^G)S^{sales}$. The mine receives operating surplus $\beta_m V$, machine invoices are $(1-\beta_m)V$, supplier manufacturing costs are $(1-\beta_m)^2 V$, and supplier profit is $\beta_m(1-\beta_m)V$. Splitting operating surplus between local wages and branch contribution proves every row of (99); the rows sum to S^{sales}. For public copying, the machine invoice equals its real resource cost $(1-\beta_m)pQ$, with no patent-profit row. Zero receipt, zero effective quality or zero essential site/headquarters input admits zero production; a formula dividing by a positive receipt is not used there.

Responsive shared headquarters. Substituting optimal local hiring leaves $C_{fr}^F H^{\varpi_H}$, with coefficient (100). The parent's common input therefore solves $\sup_{H\geq 0}\{B^F H^{\varpi_H} - w^L H\}$. Because $0<\varpi_H<1$, the unique positive solution is the stated H^*; the optimized value is Ψ. Its power of B is $1/(1-\varpi_H)>1$, proving strict convexity. If $B=0$, choose $H=0$ and $\Psi=0$. A finite entry or operating cost is subtracted after this variable-input maximum and can select closure. The contribution $\Psi(B+C;w)-\Psi(B;w)$ is the appropriate incremental value at fixed market state; assuming the parent keeps its old H would change it.

Exact company moments. Hold the firm's total ordinary labor $L^{all} = H + \sum_r L_r$ fixed. Write $v_r = (H_r^{site})^{1-\alpha_\Sigma} c_r \zeta_j r_r^{m_\beta}$. The pre-wage objective apart from $\kappa_\beta A_j p_j^{m_\beta}$ is $H^{\alpha_H}\sum_r v_r L_r^{\alpha_L}$. The local allocation is $L_r = (L^{all}-H)v_r^{k_L}/\sum_s v_s^{k_L}$. Thus $v_r^{k_L} = z_{fr}^{cap}$ and the remaining objective is $H^{\alpha_H}(L^{all}-H)^{\alpha_L}\Omega_f^{1-\alpha_L}$. The headquarters first-order condition gives

$$H_f = \frac{\alpha_H}{\alpha_\Sigma}L_f^{all}, \qquad L_{fr} = \frac{\alpha_L}{\alpha_\Sigma}L_f^{all}\frac{z_{fr}^{cap}}{\Omega_f}, \qquad U_f^{op} = \kappa_\beta C_\Sigma A_j p_j^{m_\beta}\Omega_f^{1-\alpha_L}(L_f^{all})^{\alpha_\Sigma}.$$

Maximizing over L_f^{all} yields

$$L_f^{all} = C_L (w^L)^{-k_\Sigma} A_j^{k_\Sigma} p_j^{m_\beta k_\Sigma} \Omega_f^{q_\Omega}.$$

Each site's pre-wage contribution is its capacity share of U_f^{op}. Since $Q_r = U_r^{op}/(\beta_m r_r p_j)$, adding site output gives the $\mathcal{K}_j$ moment and (102). No representative household or average site was introduced. Nonlinear headquarters sharing across sectors, indivisible branch choices, taxes charged on another base, or nonlinear selling continuations generally require additional states and different moments.

Subsistence information. To see why a mean income is insufficient, consider two households with a subsistence requirement of one each. Incomes $(1, 1)$ and $(0, 2)$ have the same total, but exact shortfalls sum to zero and one respectively. The account therefore retains household entitlements, disposable incomes and legal eligibility. A wage lower bound cannot be substituted for the shortfall of a household who is unemployed or ineligible for local wages.

E.3 Certification prices: current proofs

General homogeneous supply. South's certified-to-dirty expenditure ratio is $\omega_{NS} z^{1-\varepsilon_{NS}}$. It spends $I_S/(1 + \omega_{NS} z^{1-\varepsilon_{NS}})$ on dirty output; North spends I_N only on certified output. Their total expenditure ratio consequently implies

$$\frac{zQ_C}{Q_D} = \frac{\nu + \omega_{NS} z^{1-\varepsilon_{NS}}}{1 - \nu}.$$

For $0 \leq \nu < 1$, the supply side is strictly increasing, while the right side is decreasing, with the required opposite endpoint limits. This proves existence and uniqueness of (104). Summing actual expenditures gives $I = p_D^{1+\eta} F(z)$. Pooling gives $I = p_0^{1+\eta} F_0$. Dividing proves the three exact gross and net price tests. The case $\nu = 1$ has zero dirty final demand in this faithful snapshot and is a boundary equilibrium; it is not obtained by dividing the root by $1 - \nu = 0$.

Derivation of the mixed supply functions. Let $\alpha = \alpha_\Sigma$, $m = m_\beta$, and $k = (1 - \alpha)^{-1}$ solely within this proof. With unit net retention before certification, $\mathcal{J}_j^0 = \mathcal{K}_j^0$. Uniform certified retention $r = r_C$ changes these to $\mathcal{J}_C = r^{mk} \mathcal{J}_C^0$ and $\mathcal{K}_C = r^{mk-1} \mathcal{K}_C^0$. Labor clearing yields

$$w^L = (C_L/\bar{L}^{prod})^{1/k} p_D^m (X + b)^{1/k}, \qquad X = a(rz)^{mk}.$$

Since $C_Q/C_L^\alpha = \kappa_\beta C_\Sigma/\beta_m$, substitution gives exactly (105). Competitive dirty copying is already in b; it does not affect this calculation's exponents. This proves the homogeneity degree $m - 1 > 0$.

The logarithmic derivative of q_C with respect to z is $mk-1-\alpha mkX/(X+b) > m-1 > 0$. Dirty mobile supply is decreasing in z, and coerced supply at normalized dirty price one is constant. Thus zq_C/q_D is strictly increasing from zero to infinity. Moreover,

$$\frac{d}{dX}\{(X/r+b)(X+b)^{-\alpha}\} = (X+b)^{-\alpha-1}\{(1-\alpha)X/r + b(1/r-\alpha)\} > 0.$$

Consequently F increases continuously from $D_L b^{1-\alpha} + Z^{D,coer} < F_0$ to infinity. Its crossing of F_0 is unique. For $r = 1$, the crossing is one. For $0 < r < 1$, differentiating

$$g(r) = \frac{ar^{mk-1}+b}{(ar^{mk}+b)^{\alpha}}$$

shows a strictly positive derivative, because its sign is that of $(m-1)ar^{mk} + b[mk-1-\alpha mkr] > 0$. Hence $F(1) < F_0$. At $z = 1/r$, $F(1/r) = D_L(a/r+b)(a+b)^{-\alpha} + Z^{D,coer} > F_0$. It follows that $1 < \bar{z} < 1/r$.

The sharp demand threshold and gross certified increase. The residual in (104) is increasing in z and strictly decreasing in ν, so its root increases with ν. At $\bar{z}$ its sign is negative exactly when $\nu > \bar{\nu}$ in (106). Therefore the equilibrium root exceeds $\bar{z}$, and the dirty price falls, exactly under that condition. Notice $\bar{\nu} < 1$; a negative threshold simply makes every admissible nonnegative North share sufficient.

It remains to show the certified gross price rises under the same condition; this does not follow from the dirty-price statement alone in arbitrary supply systems. For this mixed application and any $z \geq 1$, put $v = z^{-mk} \in (0, 1]$. Then

$$\frac{F(z)}{z^m} = D_L \frac{ar^{mk-1}+bv}{(ar^{mk}+bv)^{\alpha}} + Z^{D,coer} v^{1-\alpha}.$$

The derivative with respect to v of the fraction has the sign of $(1-\alpha)bv + ar^{mk-1}(r-\alpha)$. It changes sign at most once, from negative to positive; hence its maximum on $[0, 1]$ is attained at an endpoint. At zero it is $a^{1-\alpha}r^{m-1} < (a+b)^{1-\alpha}$. At one it is $g(r) \leq g(1) = (a+b)^{1-\alpha}$, strictly so for $r < 1$. Thus $F(z)/z^m < F_0$ for $r < 1, z \geq 1$, and for $r = 1, z > 1$ as well. Since a dirty-price fall requires $z > \bar{z} \geq 1$, the certified gross quote rises. The net-receipt test is different and is not implied by this inequality.

The original independent-power benchmark. Let $a_C = (1-\phi_{NS})^{\eta} B_C$ and $a_D = B_D$. Here $q_C(z) = a_C z^{\eta}$, $q_D(z) = a_D$, so the equilibrium root is

$$(1-\nu)(a_C/a_D)z^{1+\eta} - \omega_{NS} z^{1-\varepsilon_{NS}} - \nu = 0.$$

Since $F_0 = B_C + B_D$, the dirty-price crossing is $\bar{z} = (1-\phi_{NS})^{-\eta/(1+\eta)}$. Evaluating the root residual there proves (107). Also $p_C^{1+\eta} = Iz^{1+\eta}/(a_D + a_C z^{1+\eta})$ is strictly increasing in z. At the crossing $p_D = p_0$, so the strict gross-price bound follows. The retained formula therefore uses its original independent power supplies; the mixed crossing generally differs.

E.4 Endogenous income, actual dirty rewards and growth: current proofs

Public dirty knowledge and the inherited regional cost. The installed-input procurement branch has

$$\mathcal{C}_r(q; A_r) = \frac{\kappa_r}{1 + 1/\eta} A_r^{-1/\eta} q^{1+1/\eta}, \qquad \eta > 0.$$

At an actual linear receipt $p_D > 0$, strict convexity gives $q_r = A_r \kappa_r^{-\eta} p_D^{\eta}$ and operating profit $A_r \kappa_r^{-\eta} p_D^{1+\eta}/(1+\eta)$. This is a cost-side production result. It does not require selling a monopolized patent to a coerced household. The restriction $\alpha_m = \eta/(1+\eta) = 1 - \beta_m$ gives $1 + \eta = m_\beta$ and proves (109). Sharing the design makes equipment access competitive. The inherited κ_r is a procurement-cost coefficient, including household operating rent, rather than a pure real resource cost. In the installed-input specialization, $\mathcal{C}_r = c_{ur} u_r/\alpha_m$, so real inputs are $c_{ur} u_r = \alpha_m \mathcal{C}_r$, while household operating rent is $(1 - \alpha_m)\mathcal{C}_r$. That rent remains household income; it must not be charged again as manufacturing. If a research contract assigns the temporary improvement in this owned operating value to its inventor for the production period, the success-minus-failure gain from a proportional capacity improvement is exactly the positive amount stated in the text. It is this contractually assigned claim, not an invented public-knowledge monopoly, that can support dirty research. Any continuing ownership gains after diffusion, if assigned to the same research investor, must also be included in its actual contract valuation; they cannot be suppressed by calling knowledge public.

A positive unique market solution for each active network. Set $z^p = p_C/p_D$, $s_\Sigma = m_\beta k_\Sigma - 1$. Using labor clearing to eliminate w^L, the relative CES demand condition becomes

$$\begin{aligned} A_C^{k_\Sigma} \mathcal{K}_C (z^p)^{s_\Sigma + \varepsilon_M} = {} & \kappa_M A_D^{k_\Sigma} \mathcal{K}_D \\ & + \frac{\kappa_M Z^{D,coer} \beta_m}{\kappa_\beta C_\Sigma} \left[\frac{A_C^{k_\Sigma} \mathcal{J}_C (z^p)^{m_\beta k_\Sigma} + A_D^{k_\Sigma} \mathcal{J}_D}{\bar{L}^{prod}} \right]^{\alpha_\Sigma}. \end{aligned} \tag{93}$$

Divide the right side by $(z^p)^{s_\Sigma + \varepsilon_M}$. The constant positive term is strictly decreasing from infinity to zero. For the second term, its logarithmic derivative is less than or equal to $\alpha_\Sigma m_\beta k_\Sigma - (s_\Sigma + \varepsilon_M) = -(m_\beta - 1 + \varepsilon_M) < 0$. The sum is continuous and strictly decreasing from infinity to zero; it therefore crosses the positive constant $A_C^{k_\Sigma} \mathcal{K}_C$ exactly once. CES unit cost recovers

$$p_C = \left[\gamma_M^{\varepsilon_M} + (1 - \gamma_M)^{\varepsilon_M} (z^p)^{\varepsilon_M - 1} \right]^{1/(\varepsilon_M - 1)}, \qquad p_D = p_C/z^p,$$

and labor clearing recovers the unique positive wage. The same proof allows $Z^{D,coer} = 0$. It does not solve discrete entry, government policies or research incentives: those must be consistent with this conditional market block in the equilibrium construction.

Uniform price bounds. Write $A = A_C$, $k = k_\Sigma$, $m = m_\beta$, $\alpha = \alpha_\Sigma$, $\varepsilon = \varepsilon_M$, and $z = z^p$ within this proof. All constants in the estimates are uniform over the stated compact domain of network coefficients, dirty technology and available labor. Because the mobile dirty term is bounded away from zero, (93) implies

$$A^k z^{mk-1+\varepsilon} \geq c > 0, \qquad z \geq c' A^{-k/(mk-1+\varepsilon)} = c' A^{-\beta_m \theta_\Sigma}.$$

For an upper bound, concavity of the power $\alpha \in (0,1)$ gives

$$A^k z^{mk-1+\varepsilon} \leq C\{1 + A^{\alpha k} z^{\alpha m k}\}.$$

At least one term on the right is at least half of the left. The first possibility bounds z by $CA^{-\beta_m \theta_\Sigma}$. The second bounds $Az^{m-1+\varepsilon}$ by a constant. Hence

$$z \leq C' A^{-1/(m-1+\varepsilon)} \longrightarrow 0.$$

The CES price equation now gives the certified-price limit and a positive upper and lower bound for p_C, while the lower bound for z gives $p_D = O(A^{\beta_m \theta_\Sigma})$.

Labor clearing implies

$$(w^L)^k = \frac{C_L}{\bar{L}^{prod}} \{A^k p_C^{mk} \mathcal{J}_C + A_D^k p_D^{mk} \mathcal{J}_D\}.$$

The first summand has order A^k; the second is $O(A^{k\theta_\Sigma}) = o(A^k)$. Thus $w^L = \Theta(A)$, and dirty mobile employment is $O(A^{-k(1-\theta_\Sigma)}) \to 0$. This is freely mobile employment outside the targeted coerced regions. Coerced output remains in its separate supply equation throughout the proof.

Profit and fiscal envelopes. Coerced sales are bounded by $Z^{D,coer} p_D^m = O(A^{\theta_\Sigma})$. Mobile dirty gross sales are bounded by

$$C_Q (w^L)^{-\alpha k} A_D^k p_D^{mk} \mathcal{K}_D = O(A^{k(\theta_\Sigma - \alpha)}) = O(A^{e_\Sigma}).$$

The algebraic identities

$$k(\theta_\Sigma - \alpha) = \frac{1 - \alpha\beta_m(\varepsilon_M - 1)}{\mathcal{D}_\Sigma}, \qquad e_\Sigma - \theta_\Sigma = \alpha k(\theta_\Sigma - 1) < 0$$

prove the claimed exponents. Any nonnegative private or public claim bounded by a fixed multiple of these actual flows inherits the envelope. If a dated refund or liability is unbounded in this normalization, its actual value must be added; a production envelope alone does not bound it. Certified sales have order A, and their pre-wage surplus does too. The company's ordinary labor and headquarters wage bill equals $\alpha_\Sigma \sum_r U_r^{op}$, so aggregate operating company profit before separately specified bills is $(1 - \alpha_\Sigma) \sum_r U_r^{op} = \Theta(A)$. This company claim is separate from the supplier's patent income.

Vintage multipliers and own-site royalty limits. For an independent certified supplier with quote one, the fixed vintage's quantity at a licensed site is $a_v \chi_r x_{r,a} (p^C_{r,t})^{m_\beta} \Phi_{fr,t}$. Multiplying by the margin β_m proves (112). A finite number of sites, bounded fixed capacities, bounded available ordinary labor and bounded certified receipts give a uniform upper bound for M. For a maintained licensed core with compact positive own-network coefficients, the equations for H^* and L^*, together with w^L/A_C bounded above and below, give positive lower bounds for the core's H, L, Φ. Positive retained quotes and adoption coefficients then give $M \geq \underline{M} > 0$. Only the actual surviving licensed network enters this argument. In particular, a line that is displaced has zero continuation patent receipts.

Under convergence of coefficients in Corollary I.5, the vanished mobile dirty term gives convergence of w^L/A_C to a positive constant. Then C^F_{fr}/A_C and B^F_f/A_C converge to positive constants for the active parent. Formula (100) gives a positive limit of H_f, and (98) gives a positive limit of local employment. Finally

$$\frac{T_{r,t}}{A_{C,t}} = \frac{\tau^G_r h_t \chi_r x_{r,a}}{1-\beta_m} [(1-\tau^G_r) h_t]^{m_\beta - 1} \Phi_{fr,t} \longrightarrow t_r > 0.$$

This establishes an actual local fiscal scale before public expenses. It makes no assertion about profitability after entry and catch-up bills, or about government ownership of privately held site rents.

E.5 Nonempty application and essential boundaries

A concrete parameter domain. Take $\beta_m = 1/2$, $\alpha_H = 1/10$, $\alpha_L = 1/5$, $\varepsilon_M = 2$, $\gamma_M = 1/2$, one maintained certified parent with a positive core site, and one outside dirty parent with a positive fixed site. Choose $\gamma_C = 6/5$, so the predecessor limit quote is $\bar{q}^M_C = (1/2)(11/5) = 11/10 > 1$, and require the same strict dominance on the inherited certified core. Use independent certified suppliers at quote one and competitive public dirty copying at quote $1/2$. Choose positive constant site coefficients and $\bar{L}^{prod} = 1$, bounded retained receipt fractions, $A_D = 1$, a bounded positive coerced coefficient, and $A_C = G^t_A$ with $G_A > 1$. Let the dirty installed-input elasticity be $\alpha_m = 1/2$ on this application branch. All production and price hypotheses above are satisfied, with $\theta_\Sigma = 20/27$ and $e_\Sigma = 17/27$. Ordinary free labor is allocated competitively; no dirty employment floor or migration of coerced workers is required. Establishment and research must still be funded, and the incentive theorems must implement this candidate path. The example proves nonemptiness of the production/market domain rather than assuming that those remaining conditions are automatic. The smaller step $\gamma_C = 1/5$ gives $\bar{q}^M_C = 3/5 < 1$: the predecessor then undercuts the alleged unit-quote supplier, so that step cannot instantiate the maintained unit-quote application.

Failure of unrestricted mobility without site scarcity. If $\alpha_\Sigma = 1$, the fixed site has zero exponent. The common-wage economy can then allocate all free workers away from the less productive location; k_Σ and the present moments are not defined. This recovers

the earlier mobility obstruction and shows the economic role of the productive site. For $\alpha_\Sigma < 1$ a legal site still closes when its optimized variable contribution is below its actual fixed bill. Immobility of the mine does not imply immobility of its employees.

Last-dirty-supplier boundary and price counterexample. The positive outside mobile dirty-capacity condition cannot be deleted while keeping the displayed exponent. Keep the same common-wage certified economy but let dirty supply consist of one fixed-workforce coerced technology $Q_D = Zp_D^{m-1}$, $Z > 0$, and no mobile dirty firms. Then all free ordinary labor is certified and $Q_C = CA_C p_C^{m-1}$ for a positive constant C. CES relative demand yields

$$(p_C/p_D)^{m-1+\varepsilon_M} = \kappa_M Z/(CA_C).$$

Thus dirty sales have order $A_C^{m/(m-1+\varepsilon_M)} = A_C^{1/[1+\beta_m(\varepsilon_M-1)]}$, which is still sublinear for $\varepsilon_M > 1$, but has a different derivation and exponent. The mixed theorem's *bound* remains a valid weaker bound in this particular positive constant-supply example; its proof and claimed uniform domain do not extend to arbitrary exit or shrinking dirty capacity. In particular, set $Z = A_C^{-v}$, $v > 0$. Then p_D has order $A_C^{(1+v)/(m-1+\varepsilon_M)}$, violating the mixed price bound for sufficiently large v; dirty sales have order $A_C^{[m-v(\varepsilon_M-1)]/(m-1+\varepsilon_M)}$. This example separates a failure of the price bound from a failure of the monetary bound: shrinking coerced capacity alone need not raise its total revenue. Actual rewards and proposed new capacities must be recomputed under the changed ownership and production law. When $\varepsilon_M = 1$, Cobb–Douglas expenditure shares can instead keep dirty revenues of the same order as certified revenues. Strict substitutability is essential for the sublinear monetary conclusion, even without mobility.

Unequal curvature and inherited claims. For a different installed-input elasticity, coerced sales have exponent $1 + \eta$ in p_D; the dirty envelope must be rederived. The matched-curvature condition is therefore explicit rather than hidden in the old Section 2 symbols. Likewise a public design with no inventor contract and no assigned private productive claim gives zero appropriable innovation reward, although production may benefit. Conversely, if the research investor retains a continuing owned operation whose value changes after diffusion, a purely one-period formula omitting that continuing claim is not its full success-minus-failure value.

Safe sites and physical output. Coerced output becomes exactly zero locally after successful control, whereas old contractual claims retain their dates. The legal branch may operate only on an allowed site and under the inherited physical safety conditions. An indefinitely positive-output legal mine located at a critical site with positive disturbance per unit is an explicit failure of full site safety, however high its fiscal contribution. Changes in output values do not by themselves prove declining physical disturbance, safety of every critical site, or a bound on cumulative irreversible damage. Those are separate conservation propositions.

E.6 Asymmetric innovation: ownership, values and research

E.6.1 Exclusive assignment, survival and retained vintage quality

Each sector has unit-measure design lines, matching the main production integral. Project types have masses N_i, with $\sum_{i\in j} N_i \leq 1$. An exclusive uniform lottery assigns disjoint target sets. Each particular line is targeted by type i with probability N_i. Conditional technical success has probability $h_i(s_i)$, independently of the line's inherited quality. The outcomes "not assigned", "assigned and unsuccessful", and "assigned and successful" are disjoint and exhaustive. Therefore

$$\Pr(\text{replacement}) = \sum_{i\in j} N_i h_i(s_i), \qquad \Pr(\text{survival}) = 1 - \sum_{i\in j} N_i h_i(s_i).$$

This establishes the hazard from the stipulated assignment and effort system. It does not replace the inherited exponential success function by an allocation share. The scientist resource used by type i is $N_i s_i$, whereas its mass of successful target lines is $N_i h_i(s_i)$. These are different quantities.

For a targeted line with inherited quality a, success replaces it by $(1+\gamma_i)a$, while failure leaves a unchanged. Hence its expected next quality is

$$a\left[1 + \sum_{i\in j} N_i \gamma_i h_i(s_i)\right].$$

If assignment probabilities vary with quality, the aggregate instead contains the corresponding quality-weighted moment; the unweighted mean no longer suffices. Even with quality-independent probabilities this is, by itself, an expected-value statement. For the deterministic continuum application, specify a joint allocation of project outcomes that splits every inherited quality cohort in the same success proportions. Integrating those exact cohort shares gives

$$A_{j,t+1} = A_{j,t}\left[1 + \sum_{i\in j} N_i \gamma_i h_i(s_{i,t})\right].$$

It gives the main text's $G_A = 1 + \gamma_C h_C(s^*)$ when $N_C = 1$. A proof for a finite stochastic population must use the hitting-time statement below; ordinary independent success does not turn a finite random sum into this deterministic law.

Every un-replaced patent continues with its original a_v. An increase of mean A_C reflects new patents replacing other vintages; it changes the distribution and cannot be assigned to a survivor. The aggregate mean may grow geometrically while the physical quality of each survivor remains fixed. A consortium owning every future improvement or a perpetual royalty on all frontier quality would be a different ownership model.

E.6.2 Proof of surviving-vintage valuation

Fix a public history and an alive patent at production date u. Its discounted series is

$$V_u^v = \mathbb{E}_u\left[\sum_{n=0}^{\infty} \delta_R^n I_{u+n}^v \pi_{u+n}^v\right], \qquad I_u^v = 1.$$

Absolute summability permits conditioning and interchanging expectation with the series. Split off its first term and use the tower property. On displacement the future patent value is zero; otherwise the same series starts at $u+1$. This proves

$$V_u^v = \pi_u^v + \delta_R \mathbb{E}_u[\mathbf{1}_{\{\text{survival}\}} V_{u+1}^v].$$

Iterating to a finite horizon n gives the first n discounted flows and a residual $\delta_R^n \mathbb{E}_u[I_{u+n}^v V_{u+n}^v]$. For the series value, that residual is exactly a discounted tail, whose absolute value tends to zero by summability. Conversely any value satisfying the recursion and this terminal condition equals the series after passing to the limit in the iterated identity. This proves uniqueness and excludes bubbles supported by an arbitrary nonvanishing terminal claim.

If $\pi_{u+n}^v \le \beta_m a_v \overline{\mathcal{M}}$, then

$$\mathbb{E}_u \sum_{n\ge 0} \delta_R^n I_{u+n}^v |\pi_{u+n}^v| \le \frac{\beta_m a_v \overline{\mathcal{M}}}{1-\delta_R}.$$

No condition on $\delta_R G_A$ appears: a_v is fixed. For the more general bound, assume along every relevant surviving history that the multiplier is at most $C g_M^n$, and each successive conditional survival probability is at most $1-\underline{\rho}$. Conditional iteration gives $\mathbb{E}_u I_{u+n}^v \le (1-\underline{\rho})^n$. The same comparison is geometric with ratio $\delta_R g_M(1-\underline{\rho})$. Contractual costs require their own absolute summability; the original finite settlement horizon supplies it when its individual cash flows are integrable, but does not bound their economic size uniformly over states.

Under a deterministic public path and an appropriate conditional independence condition, the expected flow is $\pi_{t+1+n}^v \prod_{k=1}^n (1-\rho_{C,t+k})$. In a stochastic public economy the correct object is the joint expectation of the survival indicator and flow. In particular, multiplying an unconditional mean profit by an unconditional survival probability need not be valid when both respond to common shocks.

A constant-multiplier check makes the distinction transparent. If $\mathcal{M}_u^v = M$ and $\rho_{C,u} = \rho$ after entry, then a project's conditional award before contractual deductions is

$$\frac{\delta_R \beta_m (1+\gamma_C) A_{C,t} M}{1-\delta_R(1-\rho)}.$$

This is a computed special case, not the general CES price path. For $\rho = 1$ it reduces to one production period; for $\rho = 0$ it is the ordinary perpetuity of a fixed vintage. Neither calculation multiplies by the project's success probability; that probability enters its ex

ante effort problem. A purported value proportional to $\sum_n \delta_R^{n+1}(1-\rho)^n G_A^n A_{C,t}$ would price free future quality and is outside the selected baseline.

E.6.3 The one-period dirty advantage and later claims

Consider success and failure of a fixed dirty project at date t on the same public continuation. Suppose its owned operating flows differ only in period $t+1$, because subsequent productive access is the same after diffusion. Subtract the two discounted cash streams. The operating difference is exactly $\delta_R \mathbb{E}_t \Delta\mathcal{O}_{D,t+1}$; later production differences cancel. Installation, collateral, subsequent net payments, refunds and residual financial claims subtract as the main text's $\mathcal{C}_D$. Therefore

$$J_D = \delta_R \mathbb{E}_t \Delta\mathcal{O}_{D,t+1} - \mathcal{C}_D.$$

A later refund is not a second production period. A remaining ownership difference that creates additional future operating rent would violate the stated one-period productive-advantage condition and must be reported as a different specification, not hidden in a perpetual dirty patent term.

Dirty technology itself need not disappear. A success operating privately at $t+1$ joins the common dirty standard after that interval; the state records both common knowledge and any undiffused prototype cohort. When dirty effort becomes zero, the finite pre-existing pipeline can still diffuse. The uniform dirty-technology bound used for divergence includes this pipeline. It is only after its clearance that no further dirty research implies an unchanged common standard.

If no owned or contracted increment exists, $\Delta\mathcal{O}_D = 0$, and incremental costs are non-negative, then $J_D \leq 0$. Zero effort is then optimal and there is no strict incentive for unsupported dirty research. Positive effort is excluded if $J_D < 0$, or if every positive effort has a positive laboratory or wage cost. When $J_D = 0$, $R_D = 0$ and $w^R = 0$, positive effort can instead be an indifferent choice. This is a limitation of immediate sharing with no private productive advantage. The baseline deliberately specifies an actual one-period owned advantage, whose size must be computed; it does not invent an additional subsidy to make dirty research positive.

E.7 Fixed-pool allocation and the exact support threshold

E.7.1 Conditional existence and competitive implementation

Fix a public equilibrium continuation, a positive normalization c_t, finite complete rewards and finite type masses $N_i > 0$. Put $P_i = \widehat{J}_i + G_i$ and $\bar{S} = S - S_V > 0$. For $P_i < 0$, $s_i = 0$

uniquely maximizes $h_i(s_i)P_i - R_i(s_i) - w^R s_i$ for any $w^R \geq 0$, since $h_i(s) > 0$ at positive effort and $R_i(s), w^R s \geq 0$. Restrict those types to zero and maximize

$$\sum_{i:P_i \geq 0} N_i\{h_i(s_i)P_i - R_i(s_i)\} \quad \text{subject to} \quad s_i \geq 0, \quad \sum_i N_i s_i \leq \bar{S}.$$

The feasible set is nonempty and compact because $s_i \leq \bar{S}/N_i$. The objective is continuous and concave. An optimum exists. Standard finite-dimensional convex optimization, or the elementary supporting-hyperplane condition for this simplex, supplies a nonnegative multiplier w^R such that each chosen s_i maximizes its concave objective minus $w^R s_i$ on $[0, \infty)$ and

$$w^R \left(\bar{S} - \sum_i N_i s_i \right) = 0.$$

For completeness, differentiable optimality is

$$h_i'(s_i)P_i - R_i'(s_i) = w^R \quad (s_i > 0), \qquad h_i'(0)P_i - R_i'(0) \leq w^R \quad (s_i = 0).$$

Concavity makes these inequalities sufficient on the unbounded individual domain as well: each tangent line bounds its objective from above. Reintroducing negative-prize types at zero preserves individual optimality. Scientists receive $c_t w^R$ and the resource account clears with complementary slackness. If $P_i > 0$, strict concavity follows from $h_i''(s) = -(\nu_i^R)^2 e^{-\nu_i^R s} < 0$, even when the laboratory cost is merely convex. If $P_i = 0$ and the laboratory cost is flat, individual indifference is possible. This is why weak support and strict marginal dominance are distinct.

No normalization of owner counts to a unit scientist pool has been used. In particular, replacing this nonlinear effort problem by a linear comparison of destination prize levels would fail to prove its marginal conditions. Existing linear-allocation Lean lemmas supply accounting or a different specification; they do not directly establish this result.

E.7.2 Proof of the support threshold

At the proposed allocation $s_C = s_C^* = \bar{S}/N_C > 0$ and $s_i = 0$ for $i \neq C$, the certified owner's marginal surplus is

$$M_C = h_C'(s_C^*)(\widehat{J}_C + G_C) - R_C'(s_C^*).$$

Necessity of individual optimality and market clearing gives $w^R = M_C \geq 0$. Every rival must have either a nonpositive total prize, in which case zero is optimal, or initial marginal surplus at most w^R. Hence necessity is exactly

$$M_C \geq \mathcal{B}, \qquad \mathcal{B} = \max\{0, \max_{i \neq C}[h_i'(0)[\widehat{J}_i + G_i]_+ - R_i'(0)]\}.$$

The positive part is harmless here because the outer maximum includes zero and $R_i'(0) \geq 0$. It does *not* replace a negative reward by zero in the owner's actual objective.

Conversely $M_C \geq \mathcal{B}$ implies $h'_C(s^*_C)(\widehat{J}_C + G_C) \geq R'_C(s^*_C) + \mathcal{B} \geq 0$. Since $h'_C(s^*_C) > 0$, the total certified prize is nonnegative. Set $w^R = M_C$. Concavity and the first-order condition make s^*_C optimal; all rivals choose zero by the initial marginal inequalities or negative-prize dominance. The full resource account clears. Dividing the inequality by the positive derivative and imposing $G_C \geq 0$ yields exactly

$$G_C \geq \max\left\{0, \frac{R'_C(s^*_C) + \mathcal{B}}{h'_C(s^*_C)} - \widehat{J}_C\right\}.$$

A strict inequality $M_C > \mathcal{B}$ makes the certified prize positive and its objective strictly concave. Any allocation sending effort to a rival or leaving it unused can improve by moving a small amount toward the certified core; the stated allocation is the unique destination allocation in this restricted research market. This does not prove general-equilibrium uniqueness.

At $R'_C(s^*_C) = \mathcal{B} = 0$ and $\widehat{J}_C + G_C = 0$, the certified owner can be indifferent and full-certified research is weakly supported. With a negative total certified prize, positive effort is strictly worse than zero even when costs vanish. Thus the raw-reward threshold remains valid in this corner. Applying a positive part inside the certified objective would alter this conclusion and define a different economic model.

A matching payment $G_C = \mu \widehat{J}_C$ with $\mu \geq 0$ implements a positive additive grant only if $\widehat{J}_C > 0$. Bounds $\widehat{J}_C \geq \underline{J}_C$ and $\widehat{J}_i \leq \overline{J}_i$ give a sufficient uniform grant by monotonicity, but their failure is not a necessary condition for subsidy need. All these arguments hold rewards fixed at the consistent continuation. Changing continuation, ownership, expected survival, financing or transfer-driven demand requires computing the rewards again.

E.7.3 Verification and absence of additional scientists

Replacing S_V by a realized provision $S_{V,t}$ changes only the available mass and $s^*_{C,t} = (S - S_{V,t})/N_C$ in the date-t argument. For uniform positive growth one needs $S_{V,t} \leq \overline{S}_V < S$. If only the inherited upper bound $S_{V,t} \leq S$ is known, verification can absorb every scientist; then no strictly positive innovation rate follows. Graduation is an eligibility change among the fixed pool and does not relax this inequality. New local research slots must also satisfy the line-assignment capacity constraint. A maintained fixed S_V must be shown sufficient for the actual accumulated verification workload; historical claims cannot be declared absent merely to preserve a desired research rate.

E.8 Reward divergence, path dependence and sustainable withdrawal

E.8.1 Proof of the actual reward envelope

Quality-independent assignment gives mean successful quality $(1 + \gamma_C)A_{C,t}$. A positive actual licensed multiplier $\underline{\mathcal{M}}$ makes the first discounted production flow at least $\delta_R \beta_m (1 +$

$\gamma_C)A_{C,t}\underline{\mathcal{M}}$. Later patent cash flows are nonnegative before separately counted contractual deductions. Subtracting the upper bound $\mathcal{C}_C \le \kappa_C A_C + K_C$ proves

$$\frac{J_C}{A_C} \ge \delta_R \beta_m (1+\gamma_C)\underline{\mathcal{M}} - \kappa_C - \frac{K_C}{A_C} = j_0 - \frac{K_C}{A_C}.$$

A bounded multiplier and absolute cost bound $|\mathcal{C}_C| \le \kappa'_C A_C + K'_C$ likewise imply

$$|J_C| \le \left[\frac{\delta_R \beta_m (1+\gamma_C)\overline{\mathcal{M}}}{1-\delta_R} + \kappa'_C\right] A_C + K'_C.$$

This establishes proportional growth for new project rewards when the lower coefficient is positive. It is fully compatible with fixed quality for every surviving old patent.

The mixed-production price bound, when its actual receipt hypotheses hold, bounds a dirty one-period owned gain by a constant times $A_C^{\theta_\Sigma}$. To extend this to the *complete* reward, success-minus-failure refunds and other financial claims must have a compatible bound. Nonnegative installation costs only decrease J_D; arbitrary growing refunds would not. The assumed complete bound $J_D \le D_0 + D_1 A_C^{\theta_\Sigma}$ therefore yields

$$[J_D/A_C]_+ \le D_0/A_C + D_1 A_C^{\theta_\Sigma - 1} \longrightarrow 0$$

when $A_C \to \infty$ and $\theta_\Sigma < 1$. No conclusion that absolute J_D falls is used. The exponent from shrinking mobile dirty employment cannot be assigned to coerced workers; their separate output-price bound is the one required here.

The following limitations are immediate counterexamples to unsupported shortcuts. A newly patented line with no licensed buyers has $\mathcal{M}^v = 0$ even if total certified production is positive. A success-contingent refund growing faster than A_C can invalidate dirty reward decline despite sublinear dirty operating sales. A disappearing outside dirty supplier can invalidate the price exponent. None is repaired by a change in notation; each requires a restricted application or a different proof.

E.8.2 Proof of the invariant-region theorem

An explicit form of the general argument is useful. Suppose for all admissible unsubsidized states with technology scale A the certified marginal surplus has lower bound $m - \epsilon_C(A)$ and the maximum competing marginal surplus, including inactivity, has upper bound $\epsilon_D(A)$, where $m > 0$, $\epsilon_C, \epsilon_D \ge 0$ are nonincreasing and tend to zero. Select A^* such that $\epsilon_C(A^*) + \epsilon_D(A^*) < m$. Then every state in

$$\mathcal{G} = \{x \in \mathcal{D} : A(x) \ge A^*\}$$

satisfies the strict zero-subsidy inequality in the support theorem. On its selected full-certified continuation, require the production/equilibrium transition F_C to remain in $\mathcal{D}$ and to satisfy $A(F_C(x)) \ge A(x)$. Consequently $F_C(x) \in \mathcal{G}$. Induction proves that the

zero-subsidy continuation remains in $\mathcal{G}$ at every finite date. The same proof applies pathwise to every admissible stochastic transition preserving these two properties.

This is a conditional invariance theorem, not an existence theorem for F_C. Its application has economic content only after F_C, its accounts and the domain conditions have been supplied. Uniform positive licensed coefficients can follow from a finite menu of site networks with a maintained active certified core, strictly positive normalized net labor, bounded and positive core adoption coefficients, and a compact tax set below confiscation, provided the actual corporate optimum and participation conditions are verified on each member. All relevant coefficient minima are then positive by continuity and compactness. If a menu contains a core closure, zero productivity, confiscatory tax or exhausted finance, this construction fails. Preserving the domain is a substantive production/finance check, not an additional institution assumed to force green research.

For varying verification, take the worst certified marginal lower bound over the compact interval of possible full-certified efforts, and require its minimum to be positive. The growth lower bound uses the smallest available certified effort. If either interval endpoint degenerates so that the margin or growth lower bound vanishes, the uniform conclusion does not follow.

E.8.3 The explicit homogeneous cutoff and finite support

With quadratic laboratory costs, $h'_C(s^*) = H_C$ and $R'_C(s^*) = r_C s^*$. The reward envelope gives

$$M_C(A) \geq m_\infty - H_C K_C/A, \qquad \mathcal{B}(A) \leq \nu_D^R(D_0/A + D_1 A^{\theta_\Sigma - 1}),$$

because $R'_D(0) = 0$. At the stated A^*, each of $H_C K_C/A$, $\nu_D^R D_0/A$ and $\nu_D^R D_1 A^{\theta_\Sigma - 1}$ is at most $m_\infty/4$. Hence

$$M_C(A) \geq 3m_\infty/4 > m_\infty/2 \geq \mathcal{B}(A) \quad \text{for all } A \geq A^*.$$

The inequalities are strict with margin at least $m_\infty/4$, even at the cutoff. No appeal to a favorable tie is needed. The cohort-balanced growth law gives $A_{C,t} = G_A^t A_{C,0}$ and $G_A > 1$, so the displayed finite integer T satisfies $A_{C,T} \geq A^*$. Before T, choose any finite nonnegative grants with $G_{C,t} > \Gamma_t$ at the proposed consistent path. They implement full-certified allocation. At and after T, the zero-grant private-green inequality implements the same allocation, and the invariant-domain argument proves continued private viability. Thus the patent values used before withdrawal can consistently anticipate the post-withdrawal allocation; the payment regime changes without assuming a change in the intended research direction.

This construction remains conditional on feasible dated finance and the market domain. Transfers or laboratory expenditure that alter the actual prices beyond that domain invalidate this proposed continuation and require recomputation. It also does not prove a universal sharp capacity ratio cutoff: dirty technology, pending diffusion, company membership and contracted claims may matter beyond a ratio of mean productivities.

The bound is explicit and sufficient, while the actual marginal condition remains necessary and sufficient at consistent values.

The numerical parameter example has $m_\infty = 0.99e^{-1} - 0.02 > 0$ and $G_A = 1 + 1.2(1 - e^{-1}) > 1$. The supplier admissibility test is $0.5(2.2) = 1.1 > 1$, so a copyable predecessor does not defeat the unit quote in this parameter example. For any finite nonnegative K_C, D_0, D_1 the displayed A^* is finite, proving the theorem's inequalities have a nonempty parameter domain. The example does not identify those constants with actual auction receipts without the receipt adapter.

E.8.4 Stochastic hitting and its financing limitation

Let $Y_t = \log A_{C,t}$ be nondecreasing. Before the target, suppose

$$\Pr(Y_{t+1} - Y_t \geq \ell \mid \mathcal{F}_t) \geq p > 0.$$

Let τ be the target hitting time and $K = \lceil [\log(A^*/A_{C,0})]_+ / \ell \rceil$. The case $K = 0$ is immediate. For $K > 0$, define augmented success indicators $I_t = \mathbf{1}_{\{\tau \leq t \text{ or } Y_{t+1} - Y_t \geq \ell\}}$. Conditional success probability is at least p both before the target and afterward, when $I_t = 1$. Conditional iteration bounds each successive waiting time for an augmented success by a geometric tail $(1-p)^n$, so it has conditional mean at most $1/p$. The time of the Kth augmented success is almost surely finite and has expectation at most K/p by successive conditioning. Either the target was already reached when an augmented success was counted, or K genuine log increments have occurred; in either case the actual hitting time is no larger. This proves the result without independence of increments and does not require any post-target growth assumption.

There need not be a deterministic finite hitting bound. Independent Bernoulli qualifying events with probability $p \in (0, 1)$ give positive probability $(1-p)^n$ of no qualifying event during any specified first n periods. Nor does finite expected hitting imply a finite worst-case cash reserve for a policy with a strictly positive cash burn on every failed round: a reserve financing n failures is exhausted on an event of positive probability. A particular success-only subsidy may admit a sharper reserve when paid successes also bound the number of funded events, but that must be derived from its specific state and payment rule. It cannot be inferred from almost-sure eventual innovation alone.

E.8.5 Continuing support and unsuccessful withdrawal

If $\widehat{J}_C \to j_\infty$, then

$$\liminf G_C^{\min} \geq \left[\frac{R_C'(s^*)}{h_C'(s^*)} - j_\infty \right]_+,$$

since $\mathcal{B} \geq 0$. A strictly negative limiting private marginal surplus therefore implies a strictly positive lower bound on the required grant. For the knife-edge example $R_C'(s^*) = h_C'(s^*) j_\infty$ and $\widehat{J}_C = j_\infty - K/A_C$, the certified marginal surplus is $-h_C'(s^*)K/A_C < 0$. Inactivity

alone prevents unsupported full-certified effort; the support threshold is at least K/A_C, strictly positive at every finite state. This is a verified obstruction despite a zero limiting subsidy.

For an explicit failure of one-date withdrawal reasoning, take $R'_C(s^*) > 0$ and a current licensed market yielding a positive certified marginal surplus. Suppose the next actual corporate equilibrium closes every site licensing the relevant new line, with no compensation claim. The next patent prize is zero, so the next certified marginal surplus equals $-R'_C(s^*) < 0$. The current comparison was correct, but the proposed private-green region was not invariant. Similarly, $S_{V,t+1} = S$ eliminates mineral scientists even if the current private reward is high. These examples show why licensed-market persistence and verification feasibility are genuine hypotheses.

E.8.6 Dated support conservation

At date $t+1$, each successful project receives the contractual amount $c_t G_{C,t}/\delta_R$. Summing over the realized successful mass $Z_{C,t}$ gives exactly $Z_{C,t} c_t G_{C,t}/\delta_R$. The conditional expected mass is $N_C h_C(s^*_C)$, but the realized cash bill is not discounted or multiplied by success a second time. For finite deterministic T, bounding $Z_{C,t} \leq N_C$ and summing proves the conservative reserve in the main text. With a segregated reserve and no other withdrawals, each cumulative payment is at most the initial reserve, so every scheduled payment is feasible. This proof supplies an amount to be financed; it does not supply the money or finance researchers' own earlier bills.

E.8.7 Formal reuse and unresolved adapters

The registered vintage-value and assignment components can support the finite-outcome hazard accounting, survival recursions and series arguments at their actual hypotheses. Existing linear destination-allocation modules do not formalize the inherited nonlinear owner-effort optimum or its KKT implementation. The new algebraic support and invariant-set lemmas establish their declared scalar or transition statements only; compilation is not manuscript acceptance or a proof of the full spatial economy. No Lean compilation is claimed by this appendix source.

The remaining application obligations are explicit: an exclusive assignment/ownership implementation for every research type; integrable full claim accounts; actual Section 3–4 receipt and licensing connections; a complete producer/network solution giving the uniform coefficient domain; verification and ordinary-labor resource feasibility; cash funding before receipts; and a joint research/market/government continuation preserving that domain. Where these are not supplied, the theorems remain valid conditional results and the corresponding full-economy claim remains unresolved. The original one-period-certified, fixed-expenditure affine cutoff belongs only to its historical comparison specification.

E.9 Mixed-model fiscal policy, finance and intervention

These results use the actual selling and ownership accounts of the revised Section 5. The common-quote tax formulas below are an explicit specialization. They do not replace the nonlinear selling continuation in Sections 2–4 or discharge its economic adapter.

E.9.1 Paid catch-up and the normalized cohort value

For an uninterrupted purchased-service schedule before graduation, induction in (127) gives

$$x_{r,a} = \min\{1, x_{r,0} + a\nu_r^{\mathrm{ad}}\}.$$

Thus the first integer age with $x_{r,a} = 1$ is H_r^{ad} in (128). A fully mature initial site gives $H_r^{\mathrm{ad}} = 0$; $x_{r,0} = 0$ is allowed but its first productive flow can be zero and its first household shortfall positive. Service at age zero is paid at the control/opening date and changes implementation for the next date. The intervention attempt was paid one date earlier. Pausing services leaves x unchanged during the pause; without funded services this induction cannot claim uninterrupted graduation. Summing private and public adoption debits gives exactly the service invoice, while the provider's ordinary resources are counted once, separately from the fixed scientist account.

Under the cohort formula's deterministic scale path, divide (126) by B_{e_r} and use $\delta_G^a B_{e_r+a}/B_{e_r} = d^a$. Absolute summability permits expectation and summation to commute. The first H_r^{ad} terms remain the actual conditional age-specific fiscal flows. The remaining normalized flow is the constant r_∞^G, so its tail is $d^{H_r^{\mathrm{ad}}} r_\infty^G \sum_{j\geq 0} d^j$. Since $d < 1$, this is exactly the second term of (129). No attempt probability is included: this is the value conditional on control, which the preceding Bellman equation already weights by success and discounts to the attempt date. If graduation time is random, or mature normalized fiscal flows vary, this deterministic finite-cohort expression is unavailable without a corresponding conditional expectation over those changes.

Cash conservation and the exact reserve test. After dates $0, \ldots, t-1$, cash is $H_t = H_0 + \sum_{s<t}(C_s^{\mathrm{cash}} - D_s^{\mathrm{cash}})$. Payment at date t requires $H_t \geq D_t^{\mathrm{cash}}$, because current credits arrive later. Rearranging all these inequalities gives (131), proving necessity and sufficiency on a specified path. A history-contingent promise is feasible only if the inequalities hold at every history where payment is promised. Settled liabilities must appear once, with matching credits to the actual recipient; a deposit changes balances but is not revenue before the stipulated settlement. This argument conserves cash and does not create it.

For the finite-reserve example, n attempts debit at least $n\kappa$. Therefore $n \leq \lfloor H/\kappa \rfloor$ without credits. With independent constant success $P \in (0,1)$, exhaustion after that many failures has probability $(1-P)^{\lfloor H/\kappa \rfloor} > 0$. More generally, any positive-probability failure history of that length suffices. An infinite reserve demand is possible despite finite

discounted expected costs. Discounted solvency, expected solvency and history-by-history cash feasibility are distinct.

Private participation and the fixed-headquarters optimum. At fixed H_f, the common-quote production coefficient satisfies $a^{\mathrm{op}}(\tau) = a^{\mathrm{op}}(0)(1-\tau)^{1/\beta_m}$. Optimizing local labor raises this coefficient to $1/(1-\alpha_L)$; hence branch operating profit is its untaxed value times $(1-\tau)^{n_L}$. Stationary discounted private benefits consequently give $J^F = V_0^F(1-\tau)^{n_L} - C^F$. For $0 \le C^F \le V_0^F$, solving $J^F \ge 0$ yields $[0, \tau^{F,\max}]$. If $C^F > V_0^F$, the set is empty. This closed form presumes that tax-independent setup and adoption costs have been gathered in C^F; a tax-dependent bill must stay in the actual optimization.

Gross sales are $S(\tau) = S(0)(1-\tau)^{n_L-1}$, since after-royalty receipts supply one factor $1-\tau$. The derivative of royalty income is

$$\frac{d}{d\tau}\{\tau S(\tau)\} = S(0)(1-\tau)^{n_L-2}(1-n_L\tau).$$

It is positive below $1/n_L$ and negative above it. Public costs that do not depend on τ do not change this maximum. Intersecting with the participation interval proves (133), provided its cash and access restrictions really leave the whole interval available. If $C^F = V_0^F$, only $\tau = 0$ is weakly feasible; if $C^F = 0$, the interior gross maximum remains below one. At $\tau = 1$, the productive interior formulas are interpreted by their zero-output limit, not by dividing through zero. A strict entry convention excludes a binding ceiling, producing a supremum unless some further slack condition holds.

Responsive headquarters and the royalty condition. Write $v = \varpi_H \in (0,1)$, $B = B^F_{f,-r} + C(\tau)$, and $C(\tau) = C(0)(1-\tau)^{n_L}$. Then

$$H_f = (vB/w^L)^{1/(1-v)}, \qquad \frac{d\log H_f}{d\tau} = -\frac{n_L\omega_r}{(1-v)(1-\tau)}.$$

The branch's optimized sales are a positive constant times $(1-\tau)^{n_L-1}H_f^v$. Thus

$$\frac{d\log[\tau S_r(\tau)]}{d\tau} = \frac{1}{\tau} - \frac{n_L - 1 + n_L v\omega_r/(1-v)}{1-\tau}.$$

Setting this expression to zero gives (134). Since $0 \le \omega_r \le 1$, the displayed bounds follow. There can be several stationary candidates; the derivative identity alone proves neither global uniqueness nor optimality after participation, public costs and financing. Under a common tax on every branch, B has the same $(1-\tau)^{n_L}$ factor, so the gross-sales exponent is $n_L/(1-v) - 1 = 1/[\beta_m(1-\alpha_H-\alpha_L)] - 1$. The unique gross maximum is the stated network rate. The fiscal jurisdiction needed to levy that uniform tax is part of the benchmark.

Corporate scope and the policy comparison. At fixed $w^L > 0$, $\Psi(B; w^L)$ is a positive multiple of $B^{1/(1-v)}$. For $C > 0$,

$$\frac{\partial}{\partial B_-}[\Psi(B_- + C) - \Psi(B_-)] = \Psi'(B_- + C) - \Psi'(B_-) > 0.$$

Hence a larger existing network increases the target branch's incremental operating return. Its optimum headquarters also rises, so target sales and royalty at each common positive tax rise. Apply these comparisons date by date under the maintained-active path and unchanged target costs; discounting preserves them when the sums are absolutely integrable. Private NPV feasibility expands. Cash feasibility need not expand merely because NPV does, which is why its preservation is a separate hypothesis. With identical or lower public costs and nested feasible tax sets, evaluating the former maximizing tax in the new set proves that optimized public value cannot fall. The stationary threshold is increasing in that value for $dP > 0$. Simultaneous changes in world wages, output prices, shared-access costs, public costs, finance or corporate shutdown remove the pointwise comparison and must be solved in equilibrium.

The complete cash-state Bellman equation. Let x collect public cash, outstanding claims, the legal-entry state and other payoff-relevant variables, and write $x^W, x^F, x^S(\tau)$ for the actual dated successors. With feasible action correspondence $\mathcal{A}_t(x)$, the exact equation is

$$\begin{aligned} V_t(x) = \max\Bigg\{ & \widetilde{b}_t(x) + \delta_G \mathbb{E}_t V_{t+1}(x^W), \\ & \sup_{\tau:\,(\mathrm{try},\tau)\in\mathcal{A}_t(x)} V_t^{\mathrm{try}}(x,\tau)\Bigg\}, \\ V_t^{\mathrm{try}}(x,\tau) = & -K_t^G(x,\tau) + \delta_G \mathbb{E}_t\Big[P_t(x,\tau)U_{t+1}^{\mathrm{ctrl}}(x^S(\tau)) \\ & \qquad + (1 - P_t(x,\tau))V_{t+1}(x^F)\Big]. \end{aligned}$$

For example, waiting receives protection while failure does not and has paid an attempt bill. Consequently x^W and x^F generally have different cash balances. An empty attempt correspondence leaves waiting. If all compared future policies are financed by the actual arrangements and these balance differences impose no further constraint or payoff difference beyond the recorded net flows, the value is independent of the redundant balance coordinate and reduces to (135). This is a domain reduction, not an assumption that the cash states coincide. If the balance itself enters a terminal payoff or a binding future choice, retain it.

E.10 Proof of Theorem J.1

Stationary normalized problem. With given P, set $A^* = \max_\tau [dPu(\tau) - k(\tau)]$. Tax-independent failure continuation makes the normalized Bellman map

$$\mathcal{F}(v) = \max\{b + dv, A^* + d(1-P)v\}.$$

It is a contraction with coefficient at most $d < 1$. Its unique fixed point is

$$v^* = \max\left\{\frac{b}{1-d}, \frac{A^*}{1-d+dP}\right\}.$$

Indeed if the first displayed term is larger, substituting it makes the wait branch binding and the other no larger; the reverse comparison makes the attempt branch binding. A strict comparison makes the corresponding action strictly optimal. Multiplication by positive denominators yields $b < b^{\text{crit}}$. For constant normalized net fiscal flow $R_t^G/B_t = r^G$, summing its geometric series gives $u = r^G/(1-d)$; it is a control-date value, so dPu already includes the success delay. No second success or discount factor is inserted. For $P = 0$ and $k \geq 0$, the attempt is weakly dominated by waiting with $b \geq 0$, strictly so when $k + b > 0$. For $d = 0$, no positive-cost intervention is optimal. With $P = 1$, the same formula remains valid with denominator one.

Convergence on the changing path. For the second assertion write

$$F_t(v) = \max\{b_t + d_t v, -k_t + d_t[P_t u_{t+1} + (1-P_t)v]\}.$$

The assumed bounds make the payoff sequence bounded, and F_t is $\bar{d}$-Lipschitz on a tail. Finite-horizon recursions with uniformly bounded terminal values therefore have a unique bounded infinite-horizon limit v_t; the effect of a terminal value m dates away is at most a constant times $\bar{d}^m$. Let

$$F(v) = \max\{dv, -k + d[Pu + (1-P)v]\}.$$

The coefficients converge, so $F_t(v^*) \to F(v^*) = v^*$, where v^* is the stationary fixed point. Iteration gives

$$|v_t - v^*| \leq \sum_{j \geq 0} \bar{d}^j |F_{t+j}(v^*) - F(v^*)|.$$

The right side tends to zero by a bounded geometric majorant. The strict condition $A = dPu - k > 0$ gives $v^* = A/(1-d+dP) > 0$. The limiting gain from attempting is

$$-k + dP(u - v^*) = A - dPv^* = \frac{1-d}{1-d+dP} A > 0.$$

Since $b_t \to 0$ and the normalized values converge, the actual gain is positive on a finite tail. Uniform bounds, convergence and a common positive margin yield a common tail date. Uniform discounted tails in the definition of U^{ctrl} justify using its actual convergent

catch-up values; no absolute bound on dirty monetary income was used. If the fiscal inequality fails, for example $b = 0$ and $k > dPu$, waiting is strictly optimal in the stationary counterexample. Relative profit divergence alone is therefore insufficient.

E.11 A regular connected policy-equilibrium theorem

This result proves a fixed point after specifying the economic interfaces; it does not assume a government equilibrium as an interface. Fix finitely many atomless cohorts, finite parent membership classes, an exact aggregation rule and a finite horizon H. A pure contingent government plan records each date's wait or attempt decision and announced tax at each own finite success/failure history. A tax can include a finite-dimensional committed schedule when that is the actual contract. The finite product $\mathcal{P}_a$ of closed tax and action sets is compact. Discrete choices have the discrete topology. Let $\mathfrak{M} = \prod_a \mathcal{P}(\mathcal{P}_a)$ be the product of probability measures on these plans, with the weak topology. It is compact and convex in the locally convex space of finite signed measures. An atomistic government's deviation does not change the world distribution $\mu \in \mathfrak{M}$, but can change its own parent's headquarters, other branches and future participation. These corporate responses must be recomputed in its payoff. A specified conditional joint law links regional plans and control outcomes within each fixed parent network. The product-mixing specialization uses independent private plan draws conditional on cohort and the common public history, together with the specified control-shock law; the fixed parent map then determines the joint network distribution. Its full distribution, rather than regional marginals alone, enters the nonlinear company moments. More general correlated laws require an enlarged distribution space and their own consistency conditions.

The regular connected class requires the following verifiable interfaces.

1. A stage-and-transition operator maps each forecast μ into the prices, wages, actual selling receipts, licensed payments, unique corporate operating/adoption responses, research allocation and cohort path. It is continuous, remains in a compact positive-price domain and uses the immutable institutions and the actual parent map. The finite tree of dates and relevant histories, with each deviating plan's affected parent reoptimized, then gives a continuous payoff $J_a(p, \mu)$. Relevant endogenous aggregate constraints are imposed by this operator; they are not replaced by independently chosen public prices.

2. The feasible-plan correspondence $\mathcal{F}_a(\mu) \subseteq \mathcal{P}_a$ is nonempty, compact-valued and continuous. It imposes every private participation, public/private cash, access and physical-implementation constraint on the appropriate histories. A zero-attempt plan is included whenever its own contractual bills are funded. Disconnected or discrete alternatives are allowed; no concavity of taxes or entry is required. The common plan space must include every admissible deviation in the claimed class, not just the convenient tax candidates.

3. The same stage response verifies the research allocation, including Section 5's strict certified-directed condition throughout its stated invariant domain. Its corporate response is an actual optimum, with an explicit tie selection wherever necessary. An always-active branch, a prescribed tax, a free patent or an added worker is not used to close the response.

These are substantive regularity and implementation hypotheses. In particular, inequalities with continuous coefficients give a closed feasible graph but do not by themselves give lower hemicontinuity. A sufficient local condition is that every feasible plan can be approximated by plans whose nontrivial inequality constraints have strict slack, uniformly on a neighborhood of the forecast, with the same discrete decisions; exact accounting equalities must first be eliminated using their unique continuous solutions. For any open neighborhood of a feasible plan, choose such a slack plan inside it. Continuity preserves its slack for nearby forecasts, proving lower hemicontinuity. Compactness and the closed graph give upper hemicontinuity. If a binding setup or exhausted reserve prevents this approximation, the condition is not established and this theorem does not apply at that boundary. Required equality-only feasible actions need their own continuous selector; slack cannot approximate an equality.

By Berge's maximum theorem, the set

$$B_a(\mu) = \arg\max_{p \in \mathcal{F}_a(\mu)} J_a(p, \mu)$$

is nonempty compact and upper hemicontinuous. Let $\mathscr{B}_a(\mu)$ be all probability measures supported on $B_a(\mu)$. This correspondence has nonempty compact convex values and a closed graph: a limit of measures supported on an upper-hemicontinuous compact maximizing set assigns zero mass to every open set whose closure is eventually disjoint from that set. The product $\mathscr{B}$ maps $\mathfrak{M}$ into itself. The Kakutani–Fan–Glicksberg fixed-point theorem gives $\mu^* \in \mathscr{B}(\mu^*)$. The operator produces the consistent market path, while every realized plan in its support is a best response and satisfies the funding and entry constraints. Mixing plans is randomization by governments, not fractional construction by a firm. Optimal continuations may be selected at null private histories without changing the induced aggregate path. The claim is a policy equilibrium for atomistic jurisdictions; it is not subgame perfection against a positive-mass coordinated market deviation.

For the infinite-horizon result use the countable product of compact finite-history action spaces. Its finite-prefix projections are compact metrizable, and so is the full product. Require continuous finite-prefix stage responses, a closed, nonempty and lower-hemicontinuous correspondence of complete feasible plans, and a bound

$$\sup_{a,p,\mu} \left| \sum_{t \geq H} \delta_G^t \mathbb{E}[\text{actual normalized-compatible monetary payoff}_{a,t}] \right| \longrightarrow 0$$

after expressing discounted monetary payoffs in common initial units. A sufficient growth bound is $|\text{flow}_t| \leq CB_t$ with $B_{t+1}/B_t \leq G$ and $\delta_G G < 1$; the analogous bounds are needed

for corporate and research continuation responses. Continuous finite-horizon payoffs then converge uniformly, hence the infinite payoff is continuous. The same fixed-point argument applies on the product plan space. Equivalently, compatible finite-horizon subsequences have a diagonal limit, and any strictly profitable finite-prefix deviation would contradict the finite-horizon optimality once its uniformly small tail is removed. Arbitrarily financed horizons do not imply a common complete feasible plan: the closed feasible-plan and cash conditions must hold for the limit itself.

What the mixed technology verifies and what remains. At positive wages and finite coefficients, the explicit local-labor and headquarters choices and Ψ are continuous. Finite-cohort adoption states and paid receipts are continuous within a fixed active-entry branch. Strict private participation, strictly feasible dated reserves and uniform private-green incentives can preserve that branch locally. These facts verify components of the regularity hypotheses. They do not supply a continuous full-auction stage adapter, endogenously selected network, or globally continuous entry selection. Nor do they make an unbounded sequence of failed attempts fundable. Those are separately identified economic or formalization obligations; the earlier fixed-price government cube theorem cannot substitute for them.

E.12 Proof of Proposition J.2

Let $A_{r,n}$ be the event that region r remains uncontrolled after n attempts following T. Conditional on its entire history and continued funded attempts, success probability is at least $P_{\min}$, so

$$\Pr(A_{r,n+1}) = \mathbb{E}[\mathbf{1}_{A_{r,n}} \Pr(\text{failure at } n+1 \mid \mathcal{H}_n)] \le q \Pr(A_{r,n}) \le q^{n+1}.$$

The events decrease; continuity of probability gives $\Pr(\cap_n A_{r,n}) = 0$. This proves almost-sure finite individual success. Integration over regions with nonnegative finite scale weights and Tonelli's theorem gives the expected mass bound. Exact continuum aggregation is a separate assumption needed to replace that expectation by the realized deterministic mass. For a finite set of m critical regions, the union bound gives probability at most mq^n that any remains uncontrolled after n common attempts; their maximum success time is finite almost surely, without independent outcomes. Absorbing safe implementation, not control alone, then gives all-site safety.

If remaining dirty intensity is bounded pathwise by $\bar{d}_r g_D^n$,

$$\mathbb{E}\mathcal{D}^D_{T+n} \le \int \bar{d}_r g_D^n \Pr(A_{r,n})\, dr \le D_T (q g_D)^n.$$

For $qg_D < 1$, the geometric series proves convergence, finite crossing of each positive tolerance, and the cumulative bound $D_T/(1 - qg_D)$. At $P_{\min} = 1$, only date T's still-dirty activity remains, consistently giving D_T. Add actual pre-transition and legal disturbance to obtain (138). If $g_D = 1$ and type-specific hazards are constant, each type's tail is at most

$\bar{d}_a M^D_{a,T}/P_a$. Bounding $M^D_{a,T} \le M^D_{a,0}$ and each of the T earlier dates by $\bar{d}_a M^D_{a,0}$ recovers the displayed constant-intensity bound. Its expectation is not a pathwise assurance when realized failures are random.

Essential boundaries and counterexamples.

1. With $P = 0$, uncontrolled mass need not fall. With finite cash and $P \in (0, 1)$, the positive-probability exhaustion history above prevents an almost-sure success conclusion unless further actual finance is supplied.

2. With exact mass $M^D_{T+n} = q^n M^D_T$ and per-site physical disturbance $g^n_D \bar{d}$, aggregate disturbance is exactly $D_T(qg_D)^n$. At $qg_D = 1$ it stays positive and its cumulative sum diverges; above one it grows. A sublinear *monetary* protection bound is not a physical intensity bound.

3. On a continuum with exact aggregation and constant $P \in (0, 1)$, uncontrolled mass is positive at every finite date. Thus individual almost-sure success and vanishing aggregate mass do not establish a common finite weakest-link outcome. On a finite set the maximum success time is almost surely finite, but is not bounded by a deterministic finite deadline with probability one when arbitrarily long failure histories remain possible.

4. If a legally operated critical pit has positive disturbance outside $\mathcal{A}^{\text{safe}}_r$, complete control and maximal royalties do not establish conservation. A noncritical admissible site must actually exist and be paid for in the zero-critical-disturbance application. Residual legal damage must otherwise enter both the safe-set test and the cumulative account.

5. Two household earnings vectors can have the same mean and different subsistence bills: at floor one, incomes $(1, 1)$ require zero, whereas $(0, 2)$ require one. The income distribution cannot be discarded without an aggregation theorem.

6. The former linear local production with a wage below the common mobile-market wage can choose zero local labor. The immobile-site technology supplies an interior local-labor solution only with positive site capacity, legal retained receipts, headquarters and access. Zero capacity, no safe legal site, an empty participation set or corporate shutdown still produces no positive local fiscal opportunity.

E.13 Preserved certification calibration

This table preserves the saved static certification comparison cited in the introduction. It is not a recalibration or numerical validation of the new asymmetric growth and mixed-site economy. The entries and their denominator are unchanged from the frozen pre-revision appendix. “Muir” is its historical continuum-menu comparator; VCG is its whole-buyer

comparator. The underlying bulk-supply, ownership and market assumptions retain their original scope.

Table 12: False-origin units per 10,000 live certified units in the saved central comparisons

Mineral	Mechanism	Before / 10,000	After / 10,000
Gold	Muir	3.469514	1.465577
Gold	VCG	3.469514	1.465577
Cobalt	Muir	55.399268	24.564506
Cobalt	VCG	55.399268	24.564505
Tantalum	Muir	56.349866	24.306086
Tantalum	VCG	56.349866	24.306089
Tin	Muir	56.193583	24.237245
Tin	VCG	56.193583	24.237248
Tungsten	Muir	54.921520	23.677200
Tungsten	VCG	54.921520	23.677202

Note: Before and after refer to late discovery 0.25 and 0.50; the denominator is live certified mass after withdrawals.

The saved exercise changes late discovery from 0.25 to 0.50 and recomputes admissions and downstream effort. Its assumptions and saved data remain in the frozen earlier appendix. Its false-origin units per live certified mass must not be confused with the working illegal share of total world mine supply, a physical extraction reduction, or a conservation outcome.

G Technical remarks on Section 3 of the main text

This section collects technical remarks moved from Section 3 of the main text. No main-text result depends on them.

G.1 Market structure of other minerals

Other minerals trade in different grades and through different institutions, as the LME's cobalt liquidity window and the IEA's account of mineral markets illustrate [London Metal Exchange, 2026, International Energy Agency, 2025].

G.2 Common-prior reading of the type law

Read as a common prior, ν is a shared distribution of beliefs about the other buyers' types, and each buyer conditions on its own information.

G.3 Information measurability of procurement

Procurement uses only the information available before bidding.

G.4 Single counting of continuation value

The continuation value counts future opportunities once and excludes receipts already counted in the same terminal transaction. Laundering and storage costs are already deducted in the gross value.

G.5 Scope of the equilibrium claim and joint deviations

Online Appendix A.2.2 establishes the 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 more structure.

G.6 Collusion incentive among coordinated buyers

Public retention and independent buyer identities play separate roles. Coordinated buyers raise a further problem, since they can suppress the rival orders that set payments. The two-buyer example in Online Appendix B.4.1 exhibits this collusion incentive, which individual truthfulness does not remove. The one-buyer, independent-uniform example in Proposition A.15 illustrates the private-seller impossibility.

G.7 Extension to procurement and capacity choices

Proposition A.16 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.

G.8 Placement margins under state-dependent sales

When whether or how many packages sell varies by state, the relevant expectation averages the floor over the packages actually sold. Realized margins can still differ across market states.

G.9 Acquisition margin versus production profit

The auction's acquisition margin is not treated as the whole profit from production.

G.10 Length of the protected policy interval

The sufficient policy interval in Corollary 2.13 ends at the earlier of two points: the assumed exposure horizon, and the outside profit gap divided by the largest per-unit-policy loss the package bound allows.

G.11 Preserved set of cost-minimizing implementations

At unchanged output, the cost-minimizing implementation correspondence $\Gamma(q)$ from Section 1 also stays the same, because the regional primitives are fixed. What is preserved is a set of implementations. Within it the procurement quote, import-price vector, household allocation and coercion choice can vary. Profit is equal across these cost-minimizing implementations at a given output and downstream value, and unchanged output by itself says nothing further about the spatial pattern of mining or its environmental damage.

G.12 Laundering share of retained receipts

The share of receipts that retained uses supply is the laundering share when laundering is the only valued retained use.

H Technical remarks on Section 4 of the main text

This section collects technical remarks moved from Section 4 of the main text. It records supporting conditions, qualifications and extensions; conditions needed by a main-text result also appear with the maintained protocol in that text.

H.1 Launch provision for existing stocks

Stocks already existing at launch have a finite multiyear grace period for admission, with duration left unspecified. This is a narrative launch provision. The formal exercises retain their stated scope, and the later market model studies current primary supply without recycling or return to that supply.

H.2 Institutional roles

A certification operator maintains the service; an authorized adjudicator determines the effect of evidence and a collector enforces the financial award.

H.3 Record permanence and financial exposure

A permanent record and finite financial exposure are separate choices.

H.4 Content of purchaser losses

The losses $L_{E,j}$, $L_{L,j}$ and $L_{U,j}$ include the contractual payments, recoveries and operational consequences actually available to the purchaser.

H.5 Insurance and recourse

Insurance and recourse also change the return to evidence: an indemnity can favor early disclosure, while full coverage of avoidable losses can remove the insured party's reason to investigate.

H.6 Generality of the finite-diversion comparison

The comparison permits finite diversions without differentiability or concavity; separately optimized batch opportunities generally cannot be added.

H.7 Private information and sequential bargaining

Private information and sequential bargaining would require additional constraints.

H.8 Shared collateral and combined roles

Shared collateral is counted once across unresolved batches and vintages. If a person performs several roles, its joint acceptance constraint replaces separate entries and its assets are counted once.

H.9 Counting at first discovery

The affected quantity and each bond are counted once at the first effective discovery.

H.10 Reviewer numbers, common failures and information sharing

Splitting unchanged total effort among more reviewers need not improve discovery. Several reviewers can also share the same source or adjudication failure; independent efforts outside that common failure do not remove it. Information sharing can save duplicated work while reducing each party's incentive to investigate.

H.11 Sufficient conditions for the projection property

The projection conditions are satisfied, for example, when entry preserves transfer efficiencies, weakly raises original collectible losses and existing reporters' reservation payments under each corresponding plan, and adds a nonnegative silence bill.

H.12 Fixed concealment targets

Maintaining a fixed concealment target can require increasing payments even when the organizer could instead abandon that target and bear bounded sanctions.

H.13 Deposit exclusivity, displaced procurement and seizures

No other slot or unresolved vintage draws on a slot's deposits. Any unpaid contractual charge for clean procurement displaced before purchase would reduce the substitution margin. Seizures go to a public collector and are not needed for the service budget.

H.14 Completeness of modeled rewards and liabilities

The stated bonds and buyer losses exhaust the modeled appropriable rewards and liabilities.

H.15 Scope of withdrawal

Withdrawal governs certified settlement; custody and civil recovery follow their own procedures.

H.16 Settlement events and economic dates

Settlement events here index updates to claims; they are distinct from the economic dates and production vintages in the later dynamic model.

H.17 Definition of contamination ratios

The contamination statistics are ratios of expected physical masses with positive denominators.

H.18 Omitted origins

Origins with zero expected relevant mass are omitted from the weights.

H.19 Participation and solvency of financial support

Due diligence requires records, monitoring, supplier adjustments and audits; shared schemes can reduce duplication. Buyers' own investigation and retained losses enter their participation decisions separately. Each participant's net payoff after costs and residual liabilities must be at least as large as its outside option. Evidence can remain available after an earlier participant's assets have disappeared; a finding may then support exclusion but little financial recovery. Retained capital principal is an asset; the opportunity cost of holding it is a resource cost. Contingent solvency requires collectible assets or committed finance adequate to the promised payments.

H.20 Historical exposure capacity

The required funding base or verification productivity includes the capacity needed for historical exposure.

H.21 Transfers in the resource account

In the resource comparison, the original invoice and damages transfers are excluded from repeated counting. Fees, damages, bribes and bond forfeitures are transfers; counting them again alongside the real services or losses they finance would duplicate costs. Commodity

production, forgery resources and the outside value of diverted stock enter the wider mineral-economy account separately.

H.22 Illustrative parameters

The illustrative parameters are distinct from the saved gold calibration.

Part IV

Complete statements of the current asymmetric model

I Market fragmentation and industry change

Certification changes access to buyers without changing the physical mineral. This section derives its initial price effect, then studies research and growth with endogenous income. Sections 2–4 retain their household, selling and certification institutions. The spatial production branch below supplies a legal alternative after control; it does not replace the coerced household economy before control.

Component	Maintained asymmetric baseline
Certified innovation	A successful inventor owns the improved machine line until displacement. Its value uses that patent's own quality vintage and the replacement risk implied by research.
Uncertified innovation	Research creates a private productive advantage for one production period; the improvement then enters shared dirty knowledge. Production and the knowledge stock persist.
Science	One fixed mass serves research and verification. Allocation is competitive; legal-site graduation changes eligibility, not population.
Legal development	A frontier parent pays for shared headquarters and local catch-up, and operates an immobile mining site with mobile workers.
Contracts	One-period appropriation does not cancel delayed settlement, deposits, refunds or old liabilities under Section 3.

The five results are the certification price shock, profitability divergence, fixed science, the support threshold, and path dependence with sustainable withdrawal. Each uses its stated market and funding conditions. In particular, a certification-date comparison at given expenditure is not an assumption that expenditure remains fixed during growth.

I.1 World trading opportunities

Gold already reaches buyers in several large markets. Table 13 shows the changing recorded shares of mainland China and the Gulf Cooperation Council. These destinations give context to the alternative trading opportunities represented in the model.

Table 13: Recorded destinations of unwrought-gold imports.

Destination	Window	Share at start (%)	Share at end (%)
Mainland China	2015–2024	25.78	22.52
Gulf Cooperation Council	2015–2023	8.88	18.09

Notes: HS 710812. Each share is the destination's recorded import value from the 'World' partner category, divided by the corresponding total across all available reporting areas returned by the archived query for that product and year. The GCC contains Bahrain, Kuwait, Oman, Qatar, Saudi Arabia and the United Arab Emirates. Its 2024 total is unavailable because the UAE observation is missing. Recorded destinations show alternative markets; they are not counts of independent buyers or effects of introducing certification. *Source:* Frozen UN Comtrade destination-import archive used in the 4 October 2026 empirical companion, with original API queries retained.

Credible eligibility changes which of these trading opportunities a producer can use. Prices and production then adjust together. To determine that response, the market must retain the regional production choices behind supply and the actual contracts under which buyers purchase it.

I.2 Actual selling receipts and an immobile productive site

The production and selling institutions of Sections 2–4 remain in force. In particular, the illegal seller earns the continuation value $W_r(q;\Omega_t)$ generated by its actual auction, retention and trading opportunities. A global price is an argument of this continuation, not an identity $W_r(q;\Omega_t) = p_{D,t}q$. The production problem is

$$\sup_{q\in\mathcal{Q}_r(\Omega_t)} \{W_r(q;\Omega_t) - \mathcal{C}_r(q;\Omega_t)\}. \tag{94}$$

For a certified operator, replace W_r by its actual certified receipts net of the current ordinary contract charges. The feasibility set retains the inherited implementation, participation, custody and location conditions. Deposits, refunds, sanctions and liabilities retain their Section 4 dates and claimants. A one-period productive advantage does not extinguish these claims. A differentiable concave selling continuation supplies a marginal receipt at an interior optimum; it does not establish a common transaction price, the usual monopoly markup, or a CES demand system.

The explicit application below uses a *common-quote branch*: the maintained selling contracts actually pay a linear marginal producer receipt. This is a specialization of (94). One exact whole-buyer subclass is available: if the inherited retention value is $H(y) = p_D y$ on the relevant stocks and every buyer's marginal package value is at most $p_D\bar{q}$, efficient VCG gives $W_r(q) = p_D q$. The existing retained-use route must actually realize this value; applying the identity to current market supply also requires funded physical delivery at the stated date, counted once. For instance, a sufficiently capacious existing laundering route with zero real channel cost and no better storage return has this property when terminal buyers pay p_D. Positive channel costs, storage dates or binding capacities retain

their actual wedges and accounts. This is a narrow derived subclass, not an equivalence for arbitrary auctions. Denote the actual gross certified quote by h_t, a public royalty by τ_r^G, and, on the royalty-only branch, the producer receipt by $p_{r,t}^C = (1 - \tau_r^G)h_t$. Other current charges are included in the receipt or charged separately, exactly once. Royalties and certification charges are different payments.

The mixed production technology combines an immobile site with paid, shared headquarters services and local mobile workers:

$$\begin{aligned} \Phi_{fr,t} &= (H_r^{site})^{1-\alpha_\Sigma} H_{f,t}^{\alpha_H} L_{fr,t}^{\alpha_L}, \\ Q_{fr,t} &= \frac{\Phi_{fr,t}^{\beta_m}}{1-\beta_m} \int_0^1 \mathcal{A}_{r,t}(m)^{\beta_m} x_{fr,t}^M(m)^{1-\beta_m}\, dm. \end{aligned} \tag{95}$$

The parameters satisfy $\alpha_H, \alpha_L > 0$, $\alpha_\Sigma = \alpha_H + \alpha_L < 1$, and $0 < \beta_m < 1$. Here H_r^{site} is fixed physical capacity, $H_{f,t}$ is an ordinary paid input shared by the parent's branches, and $L_{fr,t}$ is local ordinary employment. The symbols do not redefine the inherited operating input elasticity α_m, the regional scale w_r, or verification technology. Let $A_{r,t} = \int_0^1 \mathcal{A}_{r,t}(m)\, dm$, $m_\beta = 1/\beta_m$, and $\kappa_\beta = \beta_m/(1-\beta_m)$. For an active legal branch, $A_{r,t} = \chi_r x_{r,a} A_{C,t}$, where a is age since paid opening. Access to the frontier requires the specified ownership or license and paid local adoption; a surviving old patent retains its own vintage quality.

At a common machine quote $q_j^M > 0$ and positive net mineral receipt $p_{r,t}$, conditional machine demand and pre-wage operating surplus are

$$\begin{aligned} x_{fr,t}^M(m) &= \mathcal{A}_{r,t}(m)(p_{r,t}/q_j^M)^{m_\beta}\Phi_{fr,t}, \\ Q_{fr,t} &= \frac{A_{r,t}}{1-\beta_m}(p_{r,t}/q_j^M)^{m_\beta - 1}\Phi_{fr,t}, \\ U_{fr,t}^{op} &= a_{r,t}^{op}\Phi_{fr,t}, \qquad a_{r,t}^{op} = \kappa_\beta A_{r,t} p_{r,t}^{m_\beta}(q_j^M)^{-(m_\beta-1)}. \end{aligned} \tag{96}$$

The certified independent patent owner faces constant marginal machine cost $c_M = 1 - \beta_m$. Its unconstrained optimum is $q_C^M = 1$, but the displaced predecessor is publicly copyable at c_M. For substitutable vintages of one line, a quality step $\gamma_C > 0$ therefore imposes the sales-preserving limit quote

$$\bar{q}_C^M = (1-\beta_m)(1+\gamma_C)^{\beta_m/(1-\beta_m)}. \tag{97}$$

The maintained unit-quote application imposes $\bar{q}_C^M > 1$ on every active line, including the inherited certified core, with no stronger outside design. The leading design then wins strictly at quote one; there is no exclusive franchise suppressing predecessor copying. At $\bar{q}_C^M \le 1$, a rule assigning cost ties to the leading design would give $q_C^M = \min\{1, \bar{q}_C^M\}$; otherwise strict undercutting can yield only a supremum at the limit quote. That alternative must use its actual quote in production, company moments and patent rewards. The line profit becomes $(q_C^M - c_M)a(p/q_C^M)^{m_\beta}\Phi$, rather than the unit-quote margin. The public dirty design can instead be copied competitively at $q_D^M = 1 - \beta_m$. Thus its operating coefficient is multiplied by $\zeta_D = (1-\beta_m)^{-(m_\beta-1)}$; its public design earns no monopoly margin. Dirty research remains possible: a successful researcher or contracting operator

owns the specified one-period productive advantage, followed by diffusion, rather than a perpetual dirty patent.

Proposition I.1 (Mobile hiring and separate ownership accounts)**.** *At $w_t^L > 0$, an operating branch with positive headquarters, site capacity and producer receipt chooses the unique positive employment*

$$L_{fr,t}^{*} = \left[\frac{\alpha_L a_{r,t}^{op} (H_r^{site})^{1-\alpha_\Sigma} H_{f,t}^{\alpha_H}}{w_t^L} \right]^{1/(1-\alpha_L)}. \tag{98}$$

The local wage bill is $\alpha_L U_{fr,t}^{op}$. For a certified independent machine supplier and gross sales $S_r^{sales} = h_t Q_{fr,t}$, the exhaustive operating account is

$$\begin{array}{ll} \textit{public royalty} & \tau_r^G S_r^{sales} \\ \textit{local ordinary wages} & \alpha_L \beta_m (1-\tau_r^G) S_r^{sales} \\ \textit{branch contribution before HQ} & (1-\alpha_L)\beta_m (1-\tau_r^G) S_r^{sales} \\ \textit{independent patent profit} & \beta_m (1-\beta_m)(1-\tau_r^G) S_r^{sales} \\ \textit{real machine manufacturing} & (1-\beta_m)^2 (1-\tau_r^G) S_r^{sales}. \end{array} \tag{99}$$

The branch contribution pays headquarters, site claims and remaining private bills before yielding company profit. It is not another patent prize. For competitive public dirty machines, the patent-profit row is zero and the manufacturing row is $(1-\beta_m) p_{r,t} Q_{fr,t}$.

This result resolves the zero-hiring obstruction *conditional on an operating site*: site scarcity and diminishing marginal labor product give positive local hiring at a common wage. It does not make opening compulsory or pay fixed operating, adoption or establishment bills. A parent can still decline an unprofitable site. The physical location and safety conditions remain binding; an immobile legal mine is not safe merely because it is legal.

I.3 Company networks, common wages and income

Write $k_L = 1/(1-\alpha_L)$ and $\varpi_H = \alpha_H/(1-\alpha_L) \in (0,1)$. After choosing local labor, branch contribution is $C_{fr,t}^F H_{f,t}^{\varpi_H}$, where

$$\begin{aligned} C_{fr,t}^F &= (1-\alpha_L)\alpha_L^{\alpha_L k_L} [a_{r,t}^{op} (H_r^{site})^{1-\alpha_\Sigma}]^{k_L} (w_t^L)^{-\alpha_L k_L}, \\ B_{f,t}^F &= \sum_{r \in \mathcal{N}_{f,t}} C_{fr,t}^F, \\ H_{f,t}^{*} &= (\varpi_H B_{f,t}^F / w_t^L)^{1/(1-\varpi_H)}, \\ \Psi(B;w) &= (1-\varpi_H) B (\varpi_H B / w)^{\varpi_H/(1-\varpi_H)}. \end{aligned} \tag{100}$$

The relevant private gain from a branch is the change in Ψ for its actual parent, net of actual additional bills. At fixed wages and receipts, Ψ is strictly convex in $B > 0$, so an existing network raises the incremental operating value of the same branch. This is an IO statement at given markets; endogenous wage, price and funding changes can offset it.

Exact aggregation retains company membership. Assume a finite fixed active network at a date, sector-specific headquarters, common exponents, and $A_{r,t} = c_{r,t}A_{j,t}$. Let $r_{r,t} > 0$ be the retained fraction of the sector's actual gross producer quote $p_{j,t}$, after the charges included in its linear receipt. Define $\zeta_j = (q_j^M)^{-(m_\beta - 1)}$, $k_\Sigma = 1/(1-\alpha_\Sigma)$, $q_\Omega = (1-\alpha_L)/(1-\alpha_\Sigma)$, and

$$\begin{aligned} z_{fr,t}^{cap} &= (H_r^{site})^{(1-\alpha_\Sigma)k_L}(c_{r,t}\zeta_j)^{k_L} r_{r,t}^{m_\beta k_L}, \\ \Omega_{f,t} &= \sum_{r \in \mathcal{N}_{f,t}} z_{fr,t}^{cap}, \\ \mathcal{J}_{j,t} &= \sum_{f \in j} \Omega_{f,t}^{q_\Omega}, \\ \mathcal{K}_{j,t} &= \sum_{f \in j} \Omega_{f,t}^{q_\Omega - 1} \sum_{r \in \mathcal{N}_{f,t}} \frac{z_{fr,t}^{cap}}{r_{r,t}}. \end{aligned} \tag{101}$$

The symbol z^{cap} is a capacity weight, not Boolean governance. The dirty copying coefficient enters these moments and nowhere creates a dirty patent income. Put $C_\Sigma = (\alpha_H/\alpha_\Sigma)^{\alpha_H}(\alpha_L/\alpha_\Sigma)^{\alpha_L}$, $C_L = (\alpha_\Sigma \kappa_\beta C_\Sigma)^{k_\Sigma}$, and $C_Q = \beta_m^{-1}\alpha_\Sigma^{\alpha_\Sigma k_\Sigma}(\kappa_\beta C_\Sigma)^{k_\Sigma}$. The exact aggregate equations are

$$\begin{aligned} L_{j,t}^{all} &= C_L (w_t^L)^{-k_\Sigma} A_{j,t}^{k_\Sigma} p_{j,t}^{m_\beta k_\Sigma} \mathcal{J}_{j,t}, \\ Q_{j,t}^{mobile} &= C_Q (w_t^L)^{-\alpha_\Sigma k_\Sigma} A_{j,t}^{k_\Sigma} p_{j,t}^{m_\beta k_\Sigma - 1} \mathcal{K}_{j,t}, \\ \sum_j L_{j,t}^{all} &= \bar{L}_t^{prod} = \bar{L}_t^{free} - L_t^{impl} - L_t^{verify} - L_t^{other} > 0. \end{aligned} \tag{102}$$

Headquarters is included once in L^{all}. The deducted implementation, verification and other uses here contain only ordinary labor; scientific verification time stays in S_V and is not deducted a second time. The fixed scientific pool is a different factor account: hiring headquarters workers does not manufacture scientists, and graduation only changes eligibility within that pool. Coerced workers remain outside the freely mobile pool until control succeeds.

Company moments generally cannot be replaced by average regional technology. Even with identical sites, two separate parents give moment 2, whereas placing them under one parent gives $2^{q_\Omega} > 2$. Neither these moments nor total wages recover the distribution of household income. Subsistence payments use the actual household shortfall $[\underline{c}_h - y_{h,t}^{disp}]_+$, integrated over the supported households, with the existing eligibility and payment dates.

I.4 The certification-date price shock survives with mobile labor

The immediate pooling comparison holds inherited technology and the buyer-bloc mineral budgets $I_N \geq 0, I_S > 0$ fixed *at the certification date*, before research changes technology.

Buyer bloc b has utility $n_b + I_b \log M_b$, where n_b is numeraire consumption, North's service is $M_N(c,d) = c$, and South's service is

$$M_S(c,d) = \left[\alpha^{1/\varepsilon_{NS}} c^{(\varepsilon_{NS}-1)/\varepsilon_{NS}} + (1-\alpha)^{1/\varepsilon_{NS}} d^{(\varepsilon_{NS}-1)/\varepsilon_{NS}}\right]^{\varepsilon_{NS}/(\varepsilon_{NS}-1)}, \quad 0 < \alpha < 1. \tag{103}$$

Owned resources fund these demands and the specified bills while leaving a strictly positive numeraire remainder. Homogeneity and the first-order condition then imply dated mineral expenditure I_b. Before certification, the anonymous pool delivers a common positive clean proportion and buyers cannot select their own source mix; North and South still spend I_b because the fixed composition only multiplies their mineral service. They are not fixed world expenditure along the growth path. Set $I = I_N + I_S > 0$, $0 \le \nu = I_N/I < 1$, and retain South's CES parameters $\varepsilon_{NS} > 1$ and $\omega_{NS} = \alpha/(1-\alpha) > 0$. North buys certified output; South chooses both origins.

There is a useful exact test beyond independent power supplies. Suppose conditional supply is homogeneous of degree $\eta > 0$ in the two gross quotes and write post-certification supplies as $Q_C = p_D^{\eta} q_C(z)$, $Q_D = p_D^{\eta} q_D(z)$, where $z = p_C/p_D > 0$. Suppose $z q_C(z)/q_D(z)$ is continuous, strictly increasing from zero to infinity. Let $F(z) = z q_C(z) + q_D(z)$, and let $F_0 > 0$ denote pre-certification total supply at a unit pooling quote. Market clearing is the unique root

$$(1-\nu) z \frac{q_C(z)}{q_D(z)} - \omega_{NS} z^{1-\varepsilon_{NS}} - \nu = 0, \qquad p_D = (I/F(z))^{1/(1+\eta)}, \quad p_0 = (I/F_0)^{1/(1+\eta)}. \tag{104}$$

Thus dirty price falls exactly when $F(z) > F_0$; certified gross price rises exactly when $F(z) < z^{1+\eta} F_0$. If the certified producer keeps fraction r_C, its net receipt rises only when $F(z) < (r_C z)^{1+\eta} F_0$. These tests keep gross and net prices separate.

Theorem I.2 (Mixed-site certification shock)**.** *Consider* (102)*, positive mobile capacity in both sectors, fixed positive available ordinary labor, and coerced supply* $Z^{D,coer} p_D^{m_\beta - 1}$*,* $Z^{D,coer} \ge 0$*. At the comparison date let the only changed production receipt wedge be the uniform certified charge* $\phi_{NS} \in [0,1)$*, with* $r_C = 1 - \phi_{NS}$*; other common inherited charges must already be included in the stated benchmark. With pre-charge moments* $\mathcal{J}_j^0 = \mathcal{K}_j^0$*, set*

$$a = A_C^{k_\Sigma} \mathcal{J}_C^0, \quad b = A_D^{k_\Sigma} \mathcal{J}_D^0, \quad D_L = \frac{\kappa_\beta C_\Sigma}{\beta_m} (\bar{L}^{prod})^{\alpha_\Sigma}, \quad X = a (r_C z)^{m_\beta k_\Sigma}.$$

Then $\eta = m_\beta - 1$ *and*

$$\begin{aligned} q_C(z) &= D_L a r_C^{m_\beta k_\Sigma - 1} z^{m_\beta k_\Sigma - 1} (X+b)^{-\alpha_\Sigma}, \\ q_D(z) &= D_L b (X+b)^{-\alpha_\Sigma} + Z^{D,coer}, \\ F(z) &= D_L (X/r_C + b)(X+b)^{-\alpha_\Sigma} + Z^{D,coer}, \\ F_0 &= D_L (a+b)^{1-\alpha_\Sigma} + Z^{D,coer}. \end{aligned} \tag{105}$$

There is a unique $\bar{z}$ with $F(\bar{z}) = F_0$. It equals one when $\phi_{NS} = 0$, and lies strictly between 1 and $1/r_C$ when $\phi_{NS} > 0$. Define

$$\bar{\nu} = \frac{\bar{z}q_C(\bar{z})/q_D(\bar{z}) - \omega_{NS}\bar{z}^{1-\varepsilon_{NS}}}{1 + \bar{z}q_C(\bar{z})/q_D(\bar{z})}. \tag{106}$$

The dirty price falls if and only if $\nu > \bar{\nu}$; equality gives no dirty-price change and reversal gives a rise. Whenever the dirty price falls, the certified gross price rises. The certified net receipt still requires its separate test in (104).

The inherited closed form remains an explicit benchmark. For independent regional power supplies $B_C[(1-\phi_{NS})p_C]^\eta$ and $B_D p_D^\eta$, with the original installed-input assumptions, it is

$$p_D < p_0 \quad \Longleftrightarrow \quad (1-\nu)\frac{B_C}{B_D} < \nu + \omega_{NS}(1-\phi_{NS})^{\eta(\varepsilon_{NS}-1)/(1+\eta)}. \tag{107}$$

Under that condition $p_C/p_0 > (1-\phi_{NS})^{-\eta/(1+\eta)} \geq 1$. Its proof and supply class remain intact; its coefficient B_j is not silently substituted for the mixed company's network moments.

I.5 Endogenous growth prices and relative profit divergence

For the growth application, use a competitive final producer with a two-input CES mineral technology, elasticity $\varepsilon_M > 1$, and weight $\gamma_M \in (0,1)$. With the final good as numeraire, its unit-cost condition and relative demand are

$$\gamma_M^{\varepsilon_M} p_C^{1-\varepsilon_M} + (1-\gamma_M)^{\varepsilon_M} p_D^{1-\varepsilon_M} = 1, \qquad \frac{Q_C}{Q_D} = \kappa_M (p_C/p_D)^{-\varepsilon_M}, \quad \kappa_M = \left(\frac{\gamma_M}{1-\gamma_M}\right)^{\varepsilon_M}. \tag{108}$$

This is a distinct growth closure, not the claim that the certification-date budgets stay fixed. Gross final income is $p_C Q_C + p_D Q_D$; real machine manufacturing and other intermediate or implementation resource uses are deducted to obtain disposable resources. Wages, patent profits, company profits and public receipts come from the same sales accounts. Transfers redistribute this income and do not create resources. Existence below is conditional on feasible and funded final resource use; it does not finance research or a failed intervention automatically.

There is an explicit public-knowledge dirty supply law compatible with the inherited Selva implementation branch. Retain the installed-input specialization $q = a(A_r)u^{\alpha_m}$ and its true procurement cost; on the stated common-quote branch impose the *application condition* $\alpha_m = 1 - \beta_m$, equivalently $\eta = m_\beta - 1$. Then

$$Q_D^{coer} = Z_t^{D,coer} p_D^{m_\beta - 1}, \qquad Z_t^{D,coer} = \int_{\{r: z_{r,t}^G = 0\}} A_{r,t} \kappa_r^{-\eta} \, d\mu(r), \quad \Pi_D^{coer} = \frac{Z_t^{D,coer} p_D^{m_\beta}}{m_\beta}. \tag{109}$$

Public equipment designs affect A_r and the inherited procurement-cost coefficient κ_r; they do not charge a patent margin. Procurement includes both the real input bill at c_{ur} and

household operating rent. Its transfer component is retained as household income and is not counted as manufacturing resource use. Coerced households' labor and participation remain those of the inherited regional implementation. This is a matched-curvature subclass, not a redefinition of α_m throughout Sections 2–4. When a successful dirty prototype temporarily raises the installed capacity of an owned operation from A_r to $(1+\gamma_D)A_r$, its one-period operating-profit gain on this branch is $\gamma_D A_r \kappa_r^{-\eta} p_D^{m_\beta}/m_\beta > 0$. After diffusion the improvement remains available to production, while the exclusive productive gain ends. Actual contract costs and claims are then added at their own dates when valuing research.

Theorem I.3 (A growth application with a common wage and coerced Selva)**.** *Fix an active company network with positive finite $\mathcal{J}_j, \mathcal{K}_j$ and $\bar{L}^{prod} > 0$. Equations (102), (108) and (109) have a unique positive price ratio, price pair and ordinary wage. Along a certified-directed path suppose $A_C \to \infty$, while A_D remains in a compact positive interval; the company moments are bounded above, both sectors' mobile moments are bounded away from zero, $Z^{D,coer}$ is bounded, and $\bar{L}^{prod}$ remains in a compact positive interval. Then, with*

$$\mathcal{D}_\Sigma = 1 + \beta_m(1-\alpha_\Sigma)(\varepsilon_M - 1) > 1, \qquad \theta_\Sigma = \mathcal{D}_\Sigma^{-1} < 1, \tag{110}$$

$$\begin{aligned} p_D &= O(A_C^{\beta_m \theta_\Sigma}), \\ p_C &\longrightarrow \gamma_M^{\varepsilon_M/(\varepsilon_M-1)}, \qquad w^L = \Theta(A_C), \\ p_D Q_D^{coer} &= O(A_C^{\theta_\Sigma}), \qquad p_D Q_D^{mobile} = O(A_C^{e_\Sigma}), \\ e_\Sigma &= \frac{1-\alpha_\Sigma\beta_m(\varepsilon_M-1)}{\mathcal{D}_\Sigma} < \theta_\Sigma < 1. \end{aligned} \tag{111}$$

Certified operating opportunities have order A_C. Actual protection payments or one-period dirty private gains dominated by a bounded multiple of actual dirty sales, plus bounded dated claims, are at most $D_0 + D_1 A_C^{\theta_\Sigma}$. They may increase absolutely while declining relative to certified opportunities.

The lower bound on outside dirty *capacity* does not impose a labor floor. Its chosen employment tends to zero at the competitive wage. No coerced household migrates to prove (111). On a path with no dirty research, previously diffused dirty knowledge remains in A_D; its boundedness is consistent with accumulation before that path starts. The theorem does not presume research will remain certified: the incentive and invariance results below must establish that continuation.

Corollary I.4 (The licensed core and a local fiscal opportunity)**.** *For a fixed-vintage certified patent of quality a_v serving its actual licensed branches, its surviving flow is*

$$\pi_{m,t}^v = \beta_m a_v M_{m,t}, \qquad M_{m,t} = \sum_{r \in \mathcal{L}_{m,t}} \chi_r x_{r,a} (p_{r,t}^C)^{m_\beta} \Phi_{fr,t}. \tag{112}$$

On a finite network with bounded capacities and inputs, compact positive retained receipt and adoption coefficients, and a maintained original licensed core, $M_{m,t}$ is bounded above. If the core's normalized own company coefficients stay in a compact positive domain under

the market conditions of Theorem I.3, it is also bounded away from zero while that patent survives. A newly successful frontier vintage with quality proportional to A_C therefore has first-flow opportunity of order A_C; an old surviving vintage does not acquire that growth.

Corollary I.5 (Individual royalties need individual conditions)**.** *For a maintained-active legal branch, suppose its fixed royalty converges to a value strictly between zero and one, its site and adoption coefficients have positive limits, its own parent network and coefficients converge, and the aggregate moments and available labor converge. Then $T_{r,t}/A_{C,t} \to t_r > 0$, where $T_{r,t} = \tau_r^G h_t Q_{fr,t}$ is the actual royalty. The statement is before public catch-up, household support and service bills; a positive net public value must be established separately.*

The distinction between aggregate growth and a specific site's fiscal opportunity matters for BSC. If the company declines entry, if every outside dirty company exits, if the original licensed core closes, or if the applicable auction fails the receipt envelope, the corresponding hypothesis above fails and its conclusion cannot be imported. If legal operation is unsafe at an immutable site, neither growth nor a tax receipt satisfies the unchanged physical safe-site objective. These are boundaries of the economic result, not reasons to impose an immobile workforce or an unfunded employment subsidy.

I.6 Ownership, surviving certified patents and temporary dirty advantages

The baseline is asymmetric. Certified innovation is Schumpeterian: a successful entrant owns a machine-design line until another successful invention displaces it. Uncertified production repeats every period, but its research owner keeps a successful prototype's *productive advantage* for one production interval. The improved dirty standard then becomes shared knowledge. Dirty research is neither prohibited nor made technologically ineffective. The company owning a mining site and the investor owning a machine patent are separate claimants; licensing a frontier design does not transfer the patent to the mine.

For clarity about assignment, each sector has a unit measure of design lines $m \in [0, 1]$, as in the production integral, and $A_{j,t} = \int_0^1 a_{j,t}(m)\, dm$ is mean line quality. Type i has a fixed mass N_i of potential projects, and the eligible projects in that sector satisfy $\sum_{i \in j} N_i \leq 1$. Each project receives at most one distinct line through an exclusive lottery, independent of quality and effort, resolved after effort is chosen. Thus a project values the quality distribution of its actual target, not the entire sector. Remaining lines are unassigned. The target patent is owned by someone other than its challenger. A project that also owns other claims retains those claims in both success and failure; any actual cannibalization must instead enter its incremental value. These are ownership and assignment restrictions of the stated application, not consequences of aggregate technology alone.

Let $h_i(s) = 1 - e^{-\nu_i^R s}$. Conditional on a given public allocation, a particular certified line is replaced before the next production round with probability

$$\rho_{C,t} = \sum_{i \in C} N_i h_i(s_{i,t}) \in [0,1]. \tag{113}$$

The same scientist allocation determines both the chance of winning a new right and the threat to an incumbent. No independent exogenous displacement rate is substituted. The lottery assigns projects, not scientists: scientific effort changes h_i, and the scientist account below remains $\sum_i N_i s_i$. Graduation can make an existing project eligible for certified research; it adds neither projects beyond the stipulated assignment capacity nor scientists.

A successful project at t first produces at $t+1$. Its certified quality is $(1+\gamma_C)a_m$, where a_m is its targeted line's inherited quality. Until displacement the patent retains this vintage quality. Let I^v_{t+1+n} indicate that vintage v remains in force in production period $t+1+n$, with $I^v_{t+1} = 1$. Under the machine application its actual licensed profit is

$$\pi^v_u = \beta_m a_v \mathcal{M}^v_u, \qquad \mathcal{M}^v_u = \sum_{(f,r) \in \mathcal{L}^v_u} \chi_r x_{r,a} (p^C_{r,u})^{1/\beta_m} \Phi_{fr,u}. \tag{114}$$

Only actual licensed, admissible production enters $\mathcal{L}^v_u$. The displayed machine margin uses the admissible unit-price monopoly branch of the production application. If the displaced predecessor is competitively copyable at cost $1-\beta_m$, that branch requires $(1-\beta_m)(1+\gamma_C)^{\beta_m/(1-\beta_m)} > 1$ for strict preference for the new design at its unit quote; otherwise the actual constrained supplier margin must replace β_m throughout the reward bounds. No exclusive production franchise is assumed to remove this competitor. The expression uses the vintage a_v; replacing it by the advancing frontier would award an old inventor innovations it does not own. The company's site profit, the government's royalty and the patent owner's machine margin remain distinct accounts.

Let $J_{i,t}$ be the date-t success-minus-failure value of the project owner's actual claims, excluding current research wages, laboratory costs and the separate success subsidy. Write $\mathcal{C}_{i,t}$ for the discounted incremental installation, access, collateral, later contractual payments and residual-claim adjustment, with refunds counted negatively. The baseline rewards are

$$\begin{aligned} J_{C,t} &= \mathbb{E}_t \left[\sum_{n \geq 0} \delta_R^{n+1} I^v_{t+1+n} \pi^v_{t+1+n} \right] - \mathcal{C}_{C,t}, \\ J_{D,t} &= \delta_R \mathbb{E}_t[\Delta \mathcal{O}_{D,t+1}] - \mathcal{C}_{D,t}, \qquad \widehat{J}_{i,t} = J_{i,t}/c_t. \end{aligned} \tag{115}$$

These are *conditional success* values. Success probability is applied once, in the research choice below. The dirty operating increment $\Delta\mathcal{O}$ is the actual one-period assigned operating claim on the owned prototype or contracted production gain, net of its current production and service costs. It is not a perpetual patent rent, the sector's total sales, or a VCG bid index. Once diffusion occurs, success and failure have the same subsequent productive opportunity; this cancellation is an explicit one-period-advantage condition. All outstanding settlement differences remain in $\mathcal{C}_D$. If the same investor instead owns a

continuing operation whose post-diffusion cash flows still differ causally between success and failure, those additional gains must be valued: diffusion alone does not establish their cancellation, and that ownership arrangement is a separate comparison.

For example, the existing finite claim horizon gives the dated part of either cost account as

$$\mathcal{C}_{i,t} = \delta_R \mathbb{E}_t[\Delta K_i + \Delta B_i] + \sum_{u=2}^{T_H+1} \delta_R^u \mathbb{E}_t \Delta C_{i,t+u} - \delta_R^{T_H+1} \mathbb{E}_t \Delta V^{\text{other}}_{i,t+T_H+1}. \tag{116}$$

This preserves Section 4's deposits, refunds and liabilities. A patent displacement or dirty diffusion does not cancel a debt. An owner with no appropriable dirty productive gain and no other positive incremental claim has $J_D \leq 0$ when its incremental costs are nonnegative; immediate sharing alone supplies no research prize.

Proposition I.6 (Surviving-vintage valuation)**.** *Suppose the discounted absolute patent flows and contractual differences in* (115) *are summable. The patent's production-date value satisfies the recursion*

$$V_u^v = \pi_u^v + \delta_R \mathbb{E}_u[\mathbf{1}_{\{\text{survival into } u+1\}} V_{u+1}^v],$$

with terminal condition $\lim_{n\to\infty} \delta_R^n \mathbb{E}_u[I_{u+n}^v V_{u+n}^v] = 0$ *conditional on being alive at* u*. It is the unique value satisfying this recursion and terminal condition. If* $\mathcal{M}_u^v \leq \overline{\mathcal{M}} < \infty$*,* $\delta_R < 1$ *suffices for patent-flow summability, even when the* frontier *grows faster than* $1/\delta_R$*. More generally, a multiplier bound* Cg_M^n *and conditional survival bounded by* $(1-\underline{\rho})^n$ *suffice when* $\delta_R g_M(1-\underline{\rho}) < 1$*. The growth rate of the old patent's actual market multiplier, rather than automatic inheritance of frontier quality, is the relevant growth term.*

For a deterministic public path with conditional independence of line survival and the licensed-market multiplier, survival weights are $\prod_{k=1}^n (1-\rho_{C,t+k})$. Without that independence, the joint expectation in (115) must be retained. See Appendix E.6.

I.7 One fixed scientist pool and competitive research

The research technology and laboratory-cost specification are retained. A type-i project chooses

$$s_{i,t} \in \arg\max_{s\geq 0} \{h_i(s)[\widehat{J}_{i,t} + G_{i,t}] - R_i(s) - w_t^R s\}, \tag{117}$$

where R_i is convex, nondecreasing, differentiable, $R_i(0) = 0$, and $c_t w_t^R$ is the scientist wage. The ordinary mobile-worker wage is w_t^L; these are different factor markets. The common scale $c_t > 0$ normalizes actual research costs and rewards. Taking $c_t = A_{C,t}$ in the growth application is a declared research-cost scaling, not an assumption of constant world spending. If paid laboratories use ordinary labor, the corresponding normalized marginal cost is calculated from $w_t^L / A_{C,t}$ and can instead be bounded uniformly in the results below.

The scientist mass is fixed:

$$S_V + \sum_i N_i s_{i,t} \leq S, \qquad \bar{S} = S - S_V > 0, \qquad s_C^* = \bar{S}/N_C. \tag{118}$$

S_V is the existing paid verification provision, which must cover actual verification work, including outstanding vintages. The proof does not create verification scientists when claims persist. If that workload requires a time-varying provision $S_{V,t}$, replace s_C^* by $(S - S_{V,t})/N_C$ at each date. A uniform growth conclusion then requires $S_{V,t} \leq \overline{S}_V < S$ and uniform incentive bounds over the resulting effort interval. Graduation only changes eligible destinations within this same account.

Proposition I.7 (Competitive research with a fixed pool)**.** *At any finite complete rewards, for a finite set of project types and positive N_i, a competitive scientific allocation and a nonnegative scientist wage exist. Each owner solves (117); scientific labor clears with complementary slackness. A type with a strictly negative total prize $\widehat{J}_i + G_i$ chooses zero effort. Positive-prize objectives are strictly concave. The all-certified allocation is supported exactly by the marginal comparisons in Theorem I.8; a strict comparison excludes every rival and inactivity. It does not imply that every geographically eligible certified laboratory receives scientists.*

Conditional research-market existence is not a proof that the proposed prices, corporate network and government policy jointly form an equilibrium. The preceding production application and Section 6 must supply those connections.

I.8 Subsidies, reward divergence and a self-sustained green path

I.8.1 The funded support threshold at actual values

Evaluate the complete values at a public path consistent with the proposed full-certified allocation, its replacement hazards, ownership, prices and funding. Hold rival current payments and the anticipated continuation fixed while calculating a current additive payment. Its date-t value is $c_t G_{i,t}$; the promised payment upon success at $t+1$ is $c_t G_{i,t}/\delta_R$.

Theorem I.8 (The research support threshold)**.** *With $s_C^* = (S - S_V)/N_C$, define*

$$\begin{aligned} \mathcal{B}_t &= \max\left\{0, \max_{i \neq C}\left[h_i'(0)[\widehat{J}_{i,t} + G_{i,t}]_+ - R_i'(0)\right]\right\}, \\ \Gamma_t &= \frac{R_C'(s_C^*) + \mathcal{B}_t}{h_C'(s_C^*)} - \widehat{J}_{C,t}, \qquad G_{C,t}^{\min} = [\Gamma_t]_+. \end{aligned} \tag{119}$$

A nonnegative additive success payment weakly supports full-certified research if and only if $G_{C,t} \geq G_{C,t}^{\min}$. Strict marginal dominance requires $G_{C,t} > \Gamma_t$ and $G_{C,t} \geq 0$. Thus $\Gamma_t < 0$ permits strictly private takeover, $\Gamma_t = 0$ is a research tie, and $\Gamma_t > 0$ requires positive support at these consistent values. A positive matching subsidy can implement this payment only when the underlying private prize is positive.

The theorem preserves the original marginal-return mechanism. Its J_C is now a surviving-patent value; inserting the former one-period affine certified prize would change the model. Changing future support or its general-equilibrium effects requires recomputing J and the threshold. A reward bound gives a sufficient uniform subsidy; its failure does not prove that the actual equilibrium needs support. Appendix E.7 proves the result, including zero rewards and ties.

I.8.2 What profitability divergence does and does not establish

Proposition I.9 (Actual research-reward divergence)**.** *Consider the mixed CES application of the preceding subsection along a candidate certified-directed path. Suppose the actual licensed multiplier has positive lower bound $\underline{\mathcal{M}}$ and finite upper bound, uniformly on the admissible network domain. Line assignment is quality-independent, so a successful certified project's mean initial quality is $(1+\gamma_C)A_{C,t}$. Assume*

$$\mathcal{C}_{C,t} \le \kappa_C A_{C,t} + K_C, \quad J_{D,t} \le D_0 + D_1 A_{C,t}^{\theta_\Sigma}, \qquad \theta_\Sigma = \frac{1}{1+\beta_m(1-\alpha_\Sigma)(\varepsilon_M - 1)} < 1,$$

where all four cost/bound constants are nonnegative. The dirty bound covers the actual one-period gain and *later net claims; outstanding dirty innovations may diffuse after certified research starts, and their finite pipeline is included in the dirty-technology bound. With $c_t = A_{C,t}$,*

$$\widehat{J}_{C,t} \ge j_0 - \frac{K_C}{A_{C,t}}, \quad [\widehat{J}_{D,t}]_+ \le \frac{D_0}{A_{C,t}} + D_1 A_{C,t}^{\theta_\Sigma - 1}, \quad j_0 = \delta_R \beta_m (1+\gamma_C)\underline{\mathcal{M}} - \kappa_C. \tag{120}$$

Thus dirty research opportunities become negligible relative to the certified technology scale. When $j_0 > 0$, certified private rewards eventually exceed a positive multiple of that scale. Absolute dirty income may nevertheless increase. This is a statement about private research claims; separate company-profit and fiscal bounds remain necessary for the BSC.

The first certified production period alone yields the displayed lower bound; later surviving rents are nonnegative and strengthen it. Under the corresponding absolute claim-cost bound, bounded multipliers also yield $J_C = O(A_C)$, giving two-sided proportional growth when $j_0 > 0$. A positive sector-wide aggregate does not imply a positive market for each inventor: the licensed-share condition is essential. The dirty bound uses actual producer receipts from Sections 2–4, not an unproved identification of auction receipts with price times output.

I.8.3 An invariant private-green region

A favorable comparison on the withdrawal date is insufficient. Let $\mathcal{D}$ be a stated domain of technology distributions, licensed company networks, adoption states, verification workloads, contractual claims and funding states. A certified continuation is admissible only if its market solution and legal transitions remain in $\mathcal{D}$. This domain includes the

original certified core, which must retain an actual licensed market; neither immobile sites nor aggregate certified growth alone guarantees that condition.

Theorem I.10 (Private-green invariance)**.** *Suppose throughout a forward-invariant domain $\mathcal{D}$ that actual unsubsidized values obey (120), all dirty and other competing marginal returns have the stated vanishing upper envelope, and normalized laboratory marginal costs are uniformly bounded. Suppose the certified marginal surplus has a strictly positive limiting lower bound. Then a sufficiently large certified technology level supports full-certified research strictly at every subsequent date, provided the certified technology scale is nondecreasing and the same domain and bounds persist. The resulting region is invariant under the selected* unsubsidized *equilibrium continuation. This proves persistence of that continuation; it does not assert uniqueness of all general equilibria.*

A finite set of potential legal sites, bounded adoption and tax coefficients, a maintained active licensed certified core, positive net ordinary labor and a positive outside mobile dirty production opportunity provide a concrete coefficient domain for the CES application. One must check the actual production solution, participation, prices, old claims and funding over every permitted transition in that domain. Taking their uniform bounds is legitimate once those checks hold; calling a desired continuation "invariant" without them is not a proof.

Theorem I.11 (Inherited technology, temporary support and private takeover)**.** *Consider two homogeneous research types with $N_C = N_D = 1$, fixed $S_V < S$, $s^* = S - S_V$, and $R_i(s) = r_i s^2/2$, $r_i \geq 0$. Maintain the actual reward bounds in (120) uniformly on a forward-invariant market domain for $A_C \geq A_{\min} > 0$. Define*

$$m_\infty = \nu_C^R e^{-\nu_C^R s^*} j_0 - r_C s^* > 0, \qquad H_C = \nu_C^R e^{-\nu_C^R s^*},$$

$$A^* = \max\left\{A_{\min}, \frac{4H_C K_C}{m_\infty}, \frac{4\nu_D^R D_0}{m_\infty}, \left(\frac{4\nu_D^R D_1}{m_\infty}\right)^{1/(1-\theta_\Sigma)}\right\}. \tag{121}$$

At every $A_C \geq A^*$, zero support strictly sustains full-certified research. This explicit sufficient cutoff depends on ownership, technology, research costs and the actual dirty opportunity envelope; it is not a universal cutoff inferred from a regional average.

For the deterministic continuum application, require cohort-balanced assignment and success: each inherited quality cohort has the same assigned-success fraction $h_C(s^*)$. This yields the exact aggregate law

$$A_{C,t+1} = G_A A_{C,t}, \qquad G_A = 1 + \gamma_C h_C(s^*) > 1. \tag{122}$$

Independent individual success alone is not asserted to prove this continuum identity. Starting from $A_{C,0} \geq A_{\min}$, funded payments strictly satisfying Theorem I.8 until

$$T = \max\left\{0, \left\lceil \frac{\log(A^*/A_{C,0})}{\log G_A} \right\rceil\right\}$$

reach the invariant private-green region. Payments can then cease permanently. Every later price, survival hazard and research reward is evaluated on that unsubsidized continuation. If the initial state already exceeds the cutoff, no initial support is required. Below the sufficient cutoff, actual private takeover may still occur; use the exact threshold to decide.

The benchmark's parameter set is nonempty: with $s^* = \nu_C^R = \underline{\mathcal{M}} = 1$, $\delta_R = 0.9$, $\beta_m = 0.5$, $\gamma_C = 1.2$, $\kappa_C = 0$, and $r_C = 0.02$, the limiting lower marginal surplus is $0.99e^{-1} - 0.02 > 0$. The CES exponent can be $\theta_\Sigma = 1/2$ when $\alpha_\Sigma = 1/2$ and $\varepsilon_M = 5$. The quality step also satisfies the strict unit-quote admissibility condition. These values verify the research inequalities; a particular auction/network economy must still verify its licensed-market and claim bounds. The theorem does not present this numerical example as a solved general equilibrium.

The path dependence is economic: past certified research raises inherited A_C, while dirty research can build the rival's shared standard and raise the bound on competing opportunities. The applicable cutoff therefore depends on research history and the resulting dirty technology and contracts. Temporary support changes that history. This replaces the old fixed-world-budget, one-period-certified affine cutoff, which is retained only as a historical comparison. It does not establish that all low-certification states are trapped or that research equilibrium is globally unique.

For a finite stochastic innovation economy, replace the exact hitting date by an almost-sure statement. If $\log A_C$ never falls and, conditional on every history before the cutoff, a log increment of at least $\ell > 0$ occurs with probability at least $p > 0$, then the cutoff is reached almost surely in finite time, with expected hitting time at most $\lceil [\log(A^*/A_{C,0})]_+/\ell \rceil/p$. A finite deterministic withdrawal date does not follow. Nor does almost-sure finite hitting by itself finance an unbounded sequence of unsuccessful research rounds. The support contract and the owners' laboratory/working-capital accounts must remain feasible until the random stopping time.

Proposition I.12 (Boundaries to private takeover)**.** *Relative dirty-profit decline alone does not guarantee private green growth. If, along a candidate full-certified path, $\widehat{J}_C \to j_\infty$ and $h_C'(s^*)j_\infty - R_C'(s^*) < 0$, full-certified research eventually requires a strictly positive subsidy even if all competing dirty returns vanish. At equality the answer depends on finite-state terms. For example, if $\widehat{J}_C = j_\infty - K/A_C$ with $K > 0$ and $R_C'(s^*) = h_C'(s^*)j_\infty$, the required additive payment is at least $K/A_C > 0$ at every finite state, although it tends to zero. If the certified licensed multiplier vanishes after subsidy withdrawal or a network reorganization, the invariant-region result cannot be applied. A one-date privately favorable comparison then provides no assurance of continued green research.*

Appendix E.8 proves the cutoff, persistence, stochastic extension and boundaries. These results retain the difference between immediate private viability, finite temporary support, and continuing support.

I.9 Support, verification and dated funding

The incentive payment is an obligation, not a source of finance. If the realized successful mass at date t is $Z_{C,t} \leq N_C$, the research-support cash payment at $t+1$ is exactly

$$\frac{c_t G_{C,t}}{\delta_R} Z_{C,t}.$$

Its conditional expectation is $N_C h_C(s_C^*) c_t G_{C,t}/\delta_R$ in the homogeneous application. Success is counted once; multiplying the already success-contingent cash transfer by h_C again would underfund it. For the deterministic finite-T support programme, a segregated reserve

$$\mathscr{R}_0^R \geq \sum_{t=0}^{T-1} \frac{N_C c_t G_{C,t}}{\delta_R}$$

finances all success payments without interest or further credit, provided no other use withdraws those funds. This conservative reserve is finite when the stated finite-date rewards and grants are finite. The actual cohort-balanced programme can use the exact successful masses instead. Research owners must separately finance wages, laboratory costs and deployment before receipts arrive. Paid verification, private catch-up, government subsistence support and repeated intervention failures each retain their own cash account.

Section 6 starts when certified-directed research begins privately or through such a funded programme. Its continuation must satisfy the private-green invariant conditions or remain explicitly supported. Neither successful local control nor graduation creates extra scientists. Local dirty output, operating income and protection payments become zero upon successful control, while old contracts continue to settle. The licensed-market and funding conditions, the full auction-to-producer-receipt adapter, and a joint market/research/government equilibrium remain identifiable obligations where the application has not yet supplied them.

J Government action and the transition

Certified-directed growth changes the opportunity cost of tolerating an illegal mine. This section begins when that research allocation starts privately or is implemented by the funded support of Section I. Its continuation must satisfy the private-green invariance conditions there. Certified patents survive until displacement; dirty research can earn its one-period productive advantage before diffusion. The public comparison uses the resulting company receipts and protection payments, rather than substituting either sector's research prize for a government's income.

The legal alternative is a paid branch of a frontier parent at a specified immobile site. Ordinary workers remain mobile. Local adoption is paid, the parent pays for headquarters and licensed machines, and graduation changes eligibility within the existing scientist pool. Control, opening and successful local innovation are different events. This structure

preserves the original constrained-tax problem: the government chooses a fiscal rule that leaves private legal development worthwhile, then compares the resulting public value with protection income and the cost of obtaining control.

J.1 The conservation game and its accounts

Local governments choose attempts and announced legal taxes. Parents choose feasible establishment, headquarters and operation; research owners choose scientists; households supply ordinary labor; the buyers, suppliers and verification actors retain the institutions of Sections 2–4. A government's objective is its actual discounted fiscal and protection income, after its bills. The external conservation objective remains safety at every designated critical site in Assumption 1.6; it is not an additional payoff assigned to the local government.

The public state records Boolean control, site locations and capacities, parent membership, adoption cohorts, technology, surviving patent vintages, market and certification records, and dated public and private balances and obligations. Household earnings distributions must be retained whenever their positive-part shortfalls do not aggregate. A proposed aggregate state is sufficient only if it determines these choices, payments and transitions.

Write $z^G_{r,t} \in \{0,1\}$ for control and $i^G_{r,t} \in \{0,1\}$ for an attempt. Before control, the original coercive Selva allocation continues. A government decides before current protection is paid. Waiting receives that payment; attempting forgoes it, pays $K^G_{r,t}$, and succeeds with probability $P_{r,t}$. Success first changes production at $t+1$. In particular,

$$Q^D_{r,t} = (1 - z^G_{r,t})\widetilde{Q}^D_{r,t}, \quad R^{D,\text{new}}_{r,t} = (1 - z^G_{r,t})\widetilde{R}^{D,\text{new}}_{r,t}, \quad b^{\text{paid}}_{r,t} = (1 - z^G_{r,t})(1 - i^G_{r,t})\widetilde{b}_{r,t}. \quad (123)$$

Here $R^{D,\text{new}}$ is the receipt from current local dirty operations. Success eliminates these new flows exactly; it does not cancel gold stocks, refunds, deposits or historical claims governed by Section 4. The attempt itself does not eliminate current physical dirty activity.

Let $\mathscr{B}^G_{r,t}$ denote the actual statutorily eligible sales base, and let τ^G_r be the output royalty announced before establishment. Public receipts are $T^G_{r,t} = \tau^G_r \mathscr{B}^G_{r,t}$. The actual selling continuation determines this base and the producer's after-tax optimization. For the common-quote application of Section I.2, $\mathscr{B}^G_{r,t} = \sum_f y_{fr,t} h_t Q_{fr,t}$; nonlinear selling receipts cannot be replaced by $h_t Q$ without the stated adapter. The profit-tax share ξ_a of the earlier specification is a separate instrument, with a different base.

For the original N_r liberated households, the maintained portable entitlement has exact cost

$$S^{\text{sub}}_{r,t} = \sum_{h=1}^{N_r} [\underline{c}_{rh} - y^L_{rh,t}]_+; \qquad S^{\text{sub}}_{r,t} = N_r[\underline{c}_r - w^L_t]_+ \quad \text{under equal unit labor earnings.} \quad (124)$$

An integral replaces the sum for a household continuum. The benefit is fixed from the committed allocation, not made marginal in a household's subsequent output. Entitlement follows the original beneficiary, so local hiring is not substituted for N_r. A real-basket

guarantee uses its actual price index. Subsistence equality gives zero payment, and unequal earnings cannot generally be replaced by their mean.

Let $F_{r,t}$ collect public services and any other assigned public bills, including this government's actual contribution to funded research support where the financing contract assigns one. With its specified share $D^G_{r,t}$ of paid adoption, zero in the parent-funded baseline, the fiscal flow is

$$R^G_{r,t} = T^G_{r,t} - S^{\text{sub}}_{r,t} - F_{r,t} - D^G_{r,t}. \tag{125}$$

The parent's setup, machinery, private adoption and headquarters bills remain in its account. They are not deducted again from public royalties. Any public transfer to the parent would be an additional recorded bill. For control and funded opening at date e, define

$$U^{\text{ctrl}}_{r,e}(\tau) = \mathbb{E}_e \sum_{a \geq 0} \delta^a_G R^G_{r,e+a}(\tau), \qquad 0 < \delta_G < 1. \tag{126}$$

The age $a = 0$ is the first adoption date. The formula includes every catch-up deficit and subsequent graduation; it is not a mature-site annuity beginning on the attempt date. The baseline opens at control when funded participation is satisfied. A separate corporate-delay extension must retain support and public bills during the delay and solve the parent's entry option.

Paid catch-up and its graduation date. Write e_r for the control and opening date; an attempt at t can first produce $e_r = t + 1$. At local age $a = t - e_r$, the site's effective standard is $A_{r,t} = \chi_r x_{r,a} A_{C,t}$, with $x_{r,0} \in [0, 1]$. A purchased adoption service $e^{\text{ad}}_{r,t} \in \{0, 1\}$ advances local implementation by $\nu^{\text{ad}}_r > 0$:

$$\begin{aligned} x_{r,a+1} &= \min\{1, x_{r,a} + \nu^{\text{ad}}_r e^{\text{ad}}_{r,e_r+a}\}, \\ D^F_{r,t} + D^G_{r,t} &= c^{\text{ad}}_{r,t} e^{\text{ad}}_{r,t}, \qquad c^{\text{ad}}_{r,t} > 0. \end{aligned} \tag{127}$$

The bill is due before the service advances x, so an unfunded service gives no progress. In the baseline the parent pays the whole bill: $D^F = c^{\text{ad}} e^{\text{ad}}$, $D^G = 0$. A specified cofinancing contract can split that same bill without counting it twice. The service provider's actual ordinary-labor use belongs in L^{impl}_t, and its real inputs enter the resource account once. These implementation workers do not enlarge or consume a second scientist pool. Any verification or research use must still satisfy the original scientist account.

The maintained funded opening schedule purchases $e^{\text{ad}} = 1$ until graduation and zero afterward, with its bills included in the parent's actual participation and continuation problem. Conditional on that schedule being financed and privately implemented, graduation occurs after

$$H^{\text{ad}}_r = \left\lceil \frac{1 - x_{r,0}}{\nu^{\text{ad}}_r} \right\rceil \quad \text{periods, at date } e_r + H^{\text{ad}}_r. \tag{128}$$

The case $x_{r,0} = 1$ gives immediate graduation and no adoption bill. Before graduation the branch receives the catch-up standard; at graduation local frontier research becomes eligible within the existing project and scientist constraints. If finance fails or the actual

continuation pauses adoption, retain (127) with its realized service decisions; the displayed uninterrupted graduation date no longer follows. Merely being a frontier parent does not supply unpaid local catch-up.

There is an exact cohort formula when the monetary scale satisfies $B_{e_r+a} = G^a B_{e_r}$, $d = \delta_G G < 1$, and fiscal flows are absolutely discounted-integrable: $\sum_{a\geq 0} \delta_G^a \mathbb{E}_{e_r}|R^G_{r,e_r+a}| < \infty$. Put $r_a^G = \mathbb{E}_{e_r}[R^G_{r,e_r+a}/B_{e_r+a}]$, and suppose the conditional expected normalized mature flow is constant, $r_a^G = r_\infty^G$ for all $a \geq H_r^{\mathrm{ad}}$. Then

$$\frac{U^{\mathrm{ctrl}}_{r,e_r}}{B_{e_r}} = \sum_{a=0}^{H_r^{\mathrm{ad}}-1} d^a r_a^G + \frac{d^{H_r^{\mathrm{ad}}} r_\infty^G}{1-d}. \tag{129}$$

The first sum includes the exact age-specific subsistence shortfalls and public bills; private adoption costs remain in the company's participation and cash account. For $H_r^{\mathrm{ad}} = 0$ the sum is empty. Changing wages, support, taxes or company operation generally prevents a constant mature normalized flow, in which case the full value (126) remains necessary. The cohort formula by itself implies neither profitable opening nor positive net public value. Proof: Appendix E.9.1.

J.2 Optimal taxation subject to private opening and dated finance

The company's participation test compares its complete optimum with and without the new branch. In the responsive-headquarters application of Section I.3, its operating increment is

$$\Delta\Pi^F_{r,t}(\tau) = \Psi(B^F_{f,-r,t} + C^F_{fr,t}(\tau); w_t^L) - \Psi(B^F_{f,-r,t}; w_t^L). \tag{130}$$

Actual continuation values subtract private setup, access, adoption and other contractual bills, allow the specified operating and shutdown choices, and count the changed headquarters cost once. Denote the resulting incremental opening value by $J^{F,\mathrm{open}}_{r,e}(\tau)$. The participation condition $J^{F,\mathrm{open}} \geq 0$ is the original private-feasibility restriction applied to these revised accounts.

For each payer separately, let C_s^{cash} and D_s^{cash} be the dated credits and debits, with the period's debits due before its credits. Opening or intervention is cash-feasible precisely when its initial available reserve covers

$$\mathcal{R}^{\mathrm{req}} = \max\left\{0, \sup_{t\geq 0}\left[\sum_{s<t}(D_s^{\mathrm{cash}} - C_s^{\mathrm{cash}}) + D_t^{\mathrm{cash}}\right]\right\}. \tag{131}$$

Use every relevant history for a promised contingent contract; a finite value on the successful history alone is insufficient. Research support, private adoption, public catch-up deficits and repeated failed attempts have separate reserves. A parent balance is not public money absent an actual transfer or credit contract.

The feasible announced-tax set $\mathcal{T}_{r,e}$ therefore requires private participation, the stated opening tie rule, separate private and public cash feasibility, authorized market access and a physically admissible legal implementation. The announced rule is enforceable before sunk establishment. If this set is compact and control values are continuous, its fiscal maximizer exists:

$$\tau_{r,e}^{G,*} \in \arg\max_{\tau \in \mathcal{T}_{r,e}} U_{r,e}^{\mathrm{ctrl}}(\tau). \tag{132}$$

An empty set removes this funded-opening action. It does not justify a labor floor, compensating wage or free credit. Where attempt costs or success depend on the tax, the government instead maximizes the complete attempt payoff below.

Proposition (the inherited constrained-tax architecture). At fixed headquarters, fixed world prices and a stationary common-quote opportunity, write $n_L = 1/[\beta_m(1-\alpha_L)] > 1$. Suppose the company's opening value is exactly $V_0^F(1-\tau)^{n_L} - C^F$, with $V_0^F > 0$, $0 \le C^F \le V_0^F$, and all cash constraints leave the resulting interval available. With tax-independent public costs and opening at weak participation,

$$\tau^{F,\max} = 1 - (C^F/V_0^F)^{1/n_L}, \qquad \tau^{G,*} = \min\{\beta_m(1-\alpha_L), \tau^{F,\max}\}. \tag{133}$$

If $C^F > V_0^F$, no nonnegative royalty induces opening. A strict participation requirement can make a binding optimum only a supremum. This is the original "unconstrained fiscal optimum, capped by profitable development" logic, with its production elasticity and actual private value rederived.

Responsive headquarters is the main application. At given world wage and prices, a target branch's interior *gross-royalty* stationary point obeys

$$\tau_r^G = \frac{1}{n_L[1 + \varpi_H \omega_r/(1-\varpi_H)]}, \qquad \omega_r = \frac{C_{fr}^F}{B_f^F}, \quad \varpi_H = \frac{\alpha_H}{1-\alpha_L}. \tag{134}$$

The share is evaluated at the chosen tax; all stationary points, feasibility boundaries and endpoints must be compared. An interior point lies between $\beta_m(1-\alpha_H-\alpha_L)$ and $\beta_m(1-\alpha_L)$. A uniform tax on the entire network has gross optimum $\beta_m(1-\alpha_H-\alpha_L)$, only for an authority actually entitled to tax every branch. These gross-flow formulas do not replace maximization of the catch-up value (132). Tax-dependent public costs and private funding can move or eliminate their candidates.

A genuine net-rent tax ξ_a instead multiplies the static operating objective when its full cost base is deductible. A nonrefunded setup cost can still restrict entry. Its input neutrality is not a property of the royalty. Nor does an immobile mine make a tax promise credible: without enforceable commitment, the authority's actual post-establishment deviation payoff and the parent's continuation incentives require an additional contract analysis.

J.3 The intervention threshold with waiting and chosen taxes

The exact constrained Bellman problem includes reserve balances and dated claims in the state. Its wait and failed-attempt successors generally have different cash balances. On a domain where the actual financing arrangements make all compared continuations feasible, these balances do not constrain the payoff comparison and the reduced equation is

$$V_{r,t} = \max\Bigg\{\widetilde{b}_{r,t} + \delta_G V_{r,t+1}, \max_{\tau\in\mathcal{T}_{r,t+1}}\Big[-K^G_{r,t}(\tau) + \delta_G P_{r,t}(\tau)U^{\mathrm{ctrl}}_{r,t+1}(\tau) + \delta_G(1-P_{r,t}(\tau))V_{r,t+1}\Big]\Bigg\}. \tag{135}$$

The first alternative retains tomorrow's option to act. The second discounts next-date success once and forgoes current protection. When financing binds, the state-dependent version in Appendix E.9 must be used instead of (135).

Normalize by the endogenous monetary opportunity $B_t > 0$, not fixed world expenditure. Write $d_t = \delta_G B_{t+1}/B_t$, $b_t = \widetilde{b}_{r,t}/B_t$, $k_t = K^G_{r,t}/B_t$, and $u_t = U^{\mathrm{ctrl}}_{r,t}/B_t$. In the homogeneous application $B_t = A_{C,t}$. The values and costs below use the same feasible taxes and the same market path.

Theorem J.1 (Intervention with an optimized feasible tax)**.** *Suppose the compared policies are funded and the reduced Bellman equation applies.*

1. *In a stationary normalized economy with $0 \le d < 1$, given $P > 0$, constant $b \ge 0$, a nonempty compact feasible tax set and continuous $u(\tau), k(\tau)$, let*

$$A^* = \max_{\tau\in\mathcal{T}}\{dPu(\tau) - k(\tau)\}, \qquad b^{\mathrm{crit}} = \frac{1-d}{1-d+dP}A^*.$$

   Immediate intervention, repeated after failure, is strictly optimal exactly when $b < b^{\mathrm{crit}}$; equality is indifference. If k is tax-independent, the maximizing tax is (132). If normalized legal flow R^G_t/B_t is constant at r^G, then $u = r^G/(1-d)$.

2. *On a changing path with tax-independent K^G_t, P_t, use the optimized actual control values. Assume bounded normalized primitives, $d_t \le \bar{d} < 1$ on a tail, and*

$$d_t \to d < 1, \quad P_t \to P > 0, \quad b_t \to 0, \quad k_t \to k, \quad u_t \to u > 0, \qquad k < dPu. \tag{136}$$

   Then every remaining government of that type strictly attempts at every sufficiently late date on that funded path. The conclusion is uniform across types or public histories when the bounds, convergence and strict margin are uniform.

The first assertion is exact; the second gives an eventual sufficient condition, not a specified first intervention date. Proof: Appendix E.10.

For a tax-dependent success probability, neither part permits optimizing u alone: maximize the entire Bellman action. At $P = 0$, a positive-cost attempt cannot improve control. At $P = 1$, failure vanishes, while the one-date production delay remains. Absolute dirty protection may rise before control while b_t tends to zero; after success its local new-flow component is exactly zero by (123).

Section I.3 supplies a connected endogenous-income application when its outside dirty supply, labor, ownership and market hypotheses hold. Section I.5 additionally requires positive actual local production and the appropriate network and adoption limits to obtain a positive local fiscal scale. These are stronger requirements than aggregate profit divergence. Positive net control value also requires that services, subsistence and paid adoption leave a positive discounted surplus, and that its size clears (136). Coerced workers are not used as mobile labor to obtain the declining dirty-income ratio.

Proposition (scope can strengthen the fiscal route to control). Compare two parent networks at the same world price and wage paths, target site, legal access, catch-up costs, attempt technology and protection receipts. Suppose the larger existing network increases $B^F_{f,-r}$, the maintained-active corporate solutions apply, and its actual cash arrangements preserve every previously feasible tax. Then the new branch's incremental operating profit and, at each common positive tax, its royalty weakly increase. If public costs do not increase and the larger network's private participation set contains the former set, its optimized control value weakly increases. The stationary intervention region consequently expands when attempt costs and success are unchanged. This is a conditional spatial comparative static; it does not hold fixed wages or prices in a worldwide expansion. Its proof uses the convexity of the optimized parent value, not a transfer of foreign tax receipts to the local government.

Finite-reserve boundary. An isolated authority with reserve $H < \infty$, no replenishing credits and attempts costing at least $\kappa > 0$ can fund at most $\lfloor H/\kappa \rfloor$ attempts. If failures remain possible, a sufficiently long failure history exhausts this reserve. A positive fiscal NPV does not remove this event. Finite support for research and a finite catch-up deficit therefore do not establish finance for arbitrary repeated failures. Theorem J.1's payoff conditions and the funding condition must both be verified.

J.4 Market interaction and a connected equilibrium class

Each atomistic government takes the public market path as given; the aggregate profile must generate that path. Successful controls change site production, labor demand, headquarters inputs, royalty bases and research prizes. The parent map adds a connection between jurisdictions sharing headquarters. Remaining dirty receipts can increase through scarcity, so neither universal strategic complementarity nor monotone absolute protection income follows.

An equilibrium consists of government tax and attempt best responses, feasible privately optimal corporate operation and adoption, research allocation satisfying Section 5, household labor choice, the unchanged selling and verification equilibria, market and resource clearing, and all dated budgets. Control changes the legal production option; it does not create scientists or a public credit line. Mixing resolves genuine payoff ties; the realized branch is still opened or closed, not fractionally built.

Proposition (regular fixed-parent-map policy equilibrium). For finitely many atomless regional cohorts with a fixed parent map, a finite-horizon equilibrium exists within the following regular connected class. The actual production, selling, labor and research stage response is single-valued and continuous in the cohort policy distributions; royalty choices and contingent plans range over compact sets; the correspondence of fully funded, privately implementable plans is nonempty and continuous; and all date payoffs and state transitions are continuous. The distribution specification must determine the joint plans and control outcomes within each parent network, and each unilateral deviation must reoptimize its affected parent. The resulting best-response map must remain in this specified stage domain. Corporate participation and Section 5's invariant private-green conditions hold throughout that domain. Mixed contingent policies then admit a fixed point. The statement extends to the infinite horizon when finite-prefix responses are consistent and continuous, the correspondence of complete feasible plans is nonempty, compact-valued and continuous (including lower hemicontinuity), and all monetary continuation tails vanish uniformly after normalization. Appendix E.11 gives the fixed-point and limit argument and sufficient strict-feasibility conditions for the feasible-plan continuity assumption.

The explicit headquarters solution is continuous at positive wages and finite positive network coefficients; thus it supplies this component of the stage response. It does not prove continuity or existence of the enlarged full-auction stage, nor of endogenous parent matching, joint discretionary entry, binding-finance or noncommitment variants. Those adapters remain proof obligations. The older government-only fixed point is not a proof of those variants. The proposition is a conditional existence result, not uniqueness, a cooperative-core result or acceptance of an unverified adapter.

J.5 Safe sites, individual success and cumulative disturbance

The unchanged physical target is $z_{r,t} \in \mathcal{A}_r^{\text{safe}}$ at every designated critical site under the asserted date or limiting criterion. Here $z_{r,t}$ is the full reconstructed physical allocation of Section 2, distinct from Boolean control $z_{r,t}^G$. A legal branch must have a specified admissible site and paid access and installation. In the zero-critical-disturbance application, control closes harmful extraction at the critical pit and the legal installation uses a separately specified noncritical deposit. That deposit becomes immobile when established. Keeping a harmful mine at the same protected pit requires its actual legal damage to satisfy the existing safe set; certification cannot establish this by itself. A region without a feasible safe site is outside this conservation application.

Let $M^D_{a,t}$ be scale-weighted ungoverned mass, distinct from buyer or scientist mass. If fraction $v^G_{a,t}$ attempts and success is $P_{a,t}$, then under an explicit exact continuum aggregation law,

$$M^D_{a,t+1} = (1 - v^G_{a,t}P_{a,t})M^D_{a,t}, \qquad M^C_{a,t+1} = M^C_{a,t} + v^G_{a,t}P_{a,t}M^D_{a,t}. \tag{137}$$

Without this law the corresponding integrated conditional expectation is used; independence alone is not an asserted exact law of large numbers.

Proposition J.2 (Transition, safe sites and the damage budget)**.** *Suppose all still-ungoverned regions attempt every date from T, every such attempt is funded, conditional success probabilities are at least $P_{\min} > 0$, and successful control is absorbing. Put $q = 1 - P_{\min}$.*

1. *A given region remains uncontrolled after n attempts with probability at most q^n, and succeeds at a finite random time almost surely. Expected ungoverned mass is at most $q^n M^D_T$; the analogous deterministic bound needs the exact aggregation law. For a finite set of critical sites, if success permanently implements their respective safe allocations, all sites are safe after an almost-surely finite random date. This conclusion does not require independence across sites.*

2. *Suppose actual disturbance per still-uncontrolled site at date $T+n$, on every relevant history, is bounded by its date-T envelope times g^n_D, and let D_T be the integral of those initial envelopes. If $qg_D < 1$, expected dirty disturbance tends to zero and its cumulative tail is at most $D_T/(1-qg_D)$. Every positive tolerance for expected aggregate disturbance is crossed at a finite date. The same conclusions hold deterministically only under the stated exact physical aggregation.*

3. *Let $\mathcal{D}_{<T}$ be actual cumulative pre-T disturbance, and $\mathcal{D}^{\text{legal}}_{\geq T}$ the complete additional legal-disturbance stream, each measured in the same expected or exact sense. The irreversible-damage allowance is satisfied if*

$$\mathcal{D}_{<T} + \frac{D_T}{1 - qg_D} + \mathcal{D}^{\text{legal}}_{\geq T} \leq \mathcal{B}. \tag{138}$$

   For zero legal disturbance and uniformly bounded dirty intensity $\bar{d}_a$, the familiar sufficient bound is $\sum_a \bar{d}_a M^D_{a,0}(T + 1/P_a) \leq \mathcal{B}$. It is a constant-intensity specialization, not a bound for unrestricted growing output.

Proof: Appendix E.12.

Almost-sure individual success, aggregate decline and the weakest-link conservation outcome are separate statements. With a continuum of critical sites and $P < 1$, a positive uncontrolled mass can remain at every finite date, even though each site's success time is almost surely finite. Aggregate convergence therefore supplies no common finite all-site safety date. With growing harmful intensity, $qg_D \geq 1$ can prevent even aggregate disturbance from declining. No later fiscal or biological improvement reverses damage already counted as irreversible.

J.6 A matched comparison of pooling and certification

This comparison isolates one production date, $t+1$, at which a project undertaken at t first operates. It holds fixed inherited nonfocal technology, active sector-specific company networks, ownership, assignment and machine rights. Certified-side inventors own surviving patents in both regimes; dirty inventors own the same one-period productive advantage in both. Before certification, output origins are pooled. This does not require machine users or their patent contracts to be anonymous. Certification therefore changes output-market access, not the assumed intellectual-property regime. The later survival-weighted operating opportunities are matched; contractual differences are counted separately at their actual dates. This is a conditional comparison, not a construction of two dynamic equilibria.

Write P for pooling and C for certification. Retain the mixed-site snapshot's uniform receipt wedges and fixed active networks. Put $m = m_\beta$, $\alpha = \alpha_\Sigma \in (0,1)$, $k = (1-\alpha)^{-1}$, $d = p_D/p_0$, $z = p_C/p_D$, $r = 1 - \phi_{NS} > 0$ and $X = a(rz)^{mk}$, where $a, b > 0$ are the inherited mobile capacity moments. Let $L_P, L_C > 0$ be ordinary labor actually available for production, after verification and other uses. Conditional on the actual market quotes, labor clearing gives

$$W := \frac{w_C^L}{w_P^L} = d^m \left[\frac{X+b}{a+b}\frac{L_P}{L_C}\right]^{1/k}. \tag{139}$$

When $L_P = L_C$, the original mixed-site root determines d, z at the specified budgets. If labor provision differs, its original equal-labor threshold cannot be reused: solve clearing with each actual labor supply.

Proposition J.3 (Matched production-date owner rewards)**.** *Suppose the same sites remain active, uniform sectoral receipt changes apply to each affected parent, and fixed operating or contractual bills are accounted for separately. At the matched nonfocal production state, the certified inventor's licensed machine-profit flow and a mobile dirty owner's assigned incremental operating gain have respective multipliers*

$$M_C = (rz\,d)^{mk} W^{-\alpha k}, \qquad M_D = d^{mk} W^{-\alpha k}. \tag{140}$$

The latter covers a proportional owned prototype improvement at the same fixed network, with the same claim share, when both its success and failure operating problems have the stated homogeneous scaling. It does not cover an arbitrary endogenous change of parent, entry or claim share. For a coerced installed-capacity operation with unchanged A_r, κ_r, the output multiplier is d^{m-1} and its operating-profit and proportional prototype-gain multipliers are d^m.

If the later surviving operating-flow values coincide across the two regimes, complete date-t success-minus-failure rewards obey

$$\begin{aligned} J_C^C - J_C^P &= \delta_R(M_C - 1)\pi_{t+1}^P - (\mathcal{C}_C^C - \mathcal{C}_C^P), \\ J_D^C - J_D^P &= \delta_R(M_D - 1)\Delta\mathcal{O}_{D,t+1}^P - (\mathcal{C}_D^C - \mathcal{C}_D^P). \end{aligned} \tag{141}$$

Replace M_D by d^m for the coerced claim. Each $\mathcal{C}$ retains all incremental installation, access, collateral, later payments, refunds and residual claims. A fee already included in the production receipt is not deducted twice. If later operating opportunities differ, add their actual discounted difference to these equations; the matched formula then does not establish the sign by itself.

Proof. Labor clearing has the form $(w^L)^k = C_L L^{-1} \sum_j A_j^k p_j^{mk} \mathcal{J}_j$. The pooling sum is proportional to $p_0^{mk}(a+b)$; after certification it is proportional to $p_D^{mk}(X+b)$. Division proves (139). On a fixed sector-specific parent network, both headquarters and local employment scale by $(v^m/W)^k$ under a net-receipt ratio v. Thus the site's labor composite scales by $(v^m/W)^{\alpha k}$, and a machine margin proportional to the site's net quote to power m has multiplier $v^m(v^m/W)^{\alpha k} = v^{mk}W^{-\alpha k}$. The optimized parent operating value has that same homogeneous multiplier. Applying it separately to success and failure proves the assigned incremental-gain statement. Substituting $v = rz\,d$ or $v = d$ gives (140). On the installed-input branch, $q = A_r \kappa_r^{-(m-1)} p^{m-1}$ and $\Pi = A_r \kappa_r^{-(m-1)} p^m/m$, giving the coerced multipliers. Finally split each complete owner value into its first operating flow, later operating flows and dated contractual differences. The matched later operating terms cancel; the costs do not. This proves (141). □

Actual science and verification. Use the same fixed total scientist mass S and project masses in both regimes, but the actual funded verification provisions $S_V^x < S$. Let $s_x = (S - S_V^x)/N_C$ and use a common positive date-t normalization c for $\widehat{J}_i^x = J_i^x/c$. Let R_i^x include actual laboratory costs, including ordinary-labor costs at the appropriate wage. For given rival payments, define

$$\begin{aligned} B_x &= \max\{0, \max_{i \neq C} [h_i'(0)[\widehat{J}_i^x + G_i^x]_+ - R_i^{x\prime}(0)]\}, \\ \Gamma_x &= \frac{R_C^{x\prime}(s_x) + B_x}{h_C'(s_x)} - \widehat{J}_C^x, \qquad G_{C,x}^{\min} = [\Gamma_x]_+. \end{aligned} \tag{142}$$

The existing conditional research-market theorem gives these thresholds. At a supporting full-certified allocation its scientist wage is

$$w_x^R = h_C'(s_x)(\widehat{J}_C^x + G_C^x) - R_C^{x\prime}(s_x) \geq B_x.$$

This is not the ordinary-worker wage in (139). Current scientist wages are paid in the research objective, not deducted again from J. Thus the exact comparison is

$$\Gamma_C - \Gamma_P = \frac{R_C^{C\prime}(s_C) + B_C}{h_C'(s_C)} - \frac{R_C^{P\prime}(s_P) + B_P}{h_C'(s_P)} - (\widehat{J}_C^C - \widehat{J}_C^P). \tag{143}$$

For equal verification and laboratory costs this reduces to $(B_C - B_P)/h_C'(s_P) - \Delta\widehat{J}_C$. A lower untruncated threshold need not strictly reduce an already zero grant. More verification can reduce the effort to be supported and hence lower the grant while reducing successful innovation mass $N_C h_C(s_x)$. Report both quantities. All verification wages,

ordinary-labor uses and vintage workloads must be funded and included in their respective resource accounts.

Admissible snapshot checks and a complete-cost reversal. Take $m = 2$, $\alpha = 1/2$, $k = 2$, $a = b = D_L = Z = 1$ and equal available ordinary labor. Normalize $p_0 = 1$ by setting mineral expenditure $I = \sqrt{2} + 1$. Set $\varepsilon_{NS} = 2$. For a chosen r, z, let $F = (X/r + 1)/\sqrt{X+1} + 1$, $d = \sqrt{I/F}$, $q_C = r^3 z^3/\sqrt{X+1}$, $q_D = 1/\sqrt{X+1} + 1$ and $\nu = (zq_C/q_D - \omega_{NS}/z)/(1 + zq_C/q_D)$. These equations verify both budget and relative-demand clearing.

First, $r = 1, z = 2, \omega_{NS} = 1$ gives $\nu \simeq 0.636197$, $d \simeq 0.686469$, $W \simeq 1.373890$, $M_C \simeq 2.586145$ and $M_D \simeq 0.161634$. With equal subsequent operating opportunities and equal complete costs, the first certified flow increases while either dirty operating gain decreases. For nonnegative initial operating claims, equal science provision and equal laboratory costs, $\Delta \widehat{J}_C > 0$ and $\Delta B \leq 0$ imply $\Delta\Gamma < 0$.

Second, $r = 0.8, z = 1.14, \omega_{NS} = 0.1$ gives $\nu \simeq 0.209428$, $d \simeq 0.995997$, $W \simeq 0.912379$, $M_C \simeq 0.746167$ and $M_D \simeq 1.078591$. Dirty price falls but the mobile dirty gain increases because ordinary labor becomes cheaper. The coerced gain instead falls by the factor d^2. The inequalities are strict, so nearby admissible parameters preserve them. Neither example identifies a completed research–government equilibrium.

Third, retain the first snapshot and take $c = 1$, $\delta_R = 0.9$, an actual positive first licensed flow $\pi^P_{t+1} = p_* > 0$, and the same later operating opportunities. Choose $j > 0$ below the baseline gross operating value and baseline nonnegative costs so $J^P_C = j$. Specify an additional nonrefundable installation/access bill payable upon successful deployment whose date-t value is $0.9(M_C - 1)p_* + 2j$. Its production-date amount is that value divided by 0.9. This is a paid contract, not destroyed or double-counted collateral. Then $\Delta J_C = -2j$ despite the favorable first operating flow, whence $J^C_C = -j$. Choose nonnegative dirty incremental costs large enough that both dirty rewards are nonpositive, no rival subsidy, zero laboratory marginal costs and the same positive scientific effort capacity. Then $B_P = B_C = 0$, $\Gamma_P = -j$, $\Gamma_C = j$: the required additive grant rises from zero to j. A grant above j strictly supports certified effort in the second regime. This is a contract-cost reversal, not a claim that actual certification necessarily charges this amount or that this amount solves a fraud contract.

For feasibility choose a positive funded verification provision in both regimes, a positive productive labor remainder, a strict admissible machine quality step, and finite outside-good endowments covering all stated finite bills and a positive numeraire remainder. The sunk/contingent bill does not change the conditional intensive-margin supply when the same sites stay active. Participation in research with a negative unsubsidized prize is not assumed: it requires the stated grant or else zero effort. If the extra access bill requires real services, supply those inputs and account for them; a transfer alone is not a new real resource. Fraud or discovery differences must instead be calculated from an actual feasible

Section 4 contract with its exclusive bonds, reporting incentives, funding and outstanding liabilities. No arbitrary fraud rate is inferred from these examples.

Limits. Matching the nonfocal production state while allowing different research effort is a conditional reward experiment. Different aggregate science allocations generally change next-period technology; these comparisons do not construct a common fully endogenous transition. Endogenous later patent replacement, changing licensed networks, fraud, entry, government control and finance require their own consistent continuation. The comparison establishes neither a universal favorable certification sign nor that output certification creates patent enforceability.

K Further results under the current model

The results below retain the current asymmetric ownership, production, resource, timing and funding conditions. They strengthen the matched-state comparison and the sufficient private-continuation certificate; they do not compare two complete dynamic equilibrium paths. The North–South snapshot and the CES growth economy are distinct demand specifications. The snapshot uses $m = 1/\beta_m$, $\alpha = \alpha_\Sigma$, and retained certified receipt $r = 1 - \phi_{NS}$. Certified patents survive until displacement; the dirty prototype advantage lasts one production interval in both regimes. The unit-quote machine margin requires

$$(1 - \beta_m)(1 + \gamma_C)^{\beta_m/(1-\beta_m)} > 1.$$

Outside this branch the actual constrained margin replaces the displayed unit-quote margin. A positive private value does not establish that research, verification, deployment or repeated public attempts are funded.

K.1 Certification, spending and matched operating rewards

This section derives sharper statements within the current model's existing mixed-site certification-date comparison. It does not substitute the earlier independent-power-supply economy or its one-period certified ownership rule. The production date, nonfocal inherited technologies, active sector-specific company networks, assignment and ownership rights are matched as in the current comparison. In particular, certified patents remain owned until displacement and the dirty prototype advantage lasts one production interval in both regimes. The formulas below concern first operating flows; all later operating differences and actual contractual and research costs must still enter complete owner values.

K.1.1 The existing mixed-site market

Retain the certification-date assumptions already imposed in the current mixed-site theorem: positive ordinary labor available for production is the same in both regimes; total dated mineral spending is $I > 0$; mobile capacity is positive in both sectors; networks remain active; and the certified receipt charge is the only changed uniform receipt wedge. Write

$$m > 1, \qquad 0 < \alpha < 1, \qquad k = (1 - \alpha)^{-1}, \qquad a, b, D_L > 0, \qquad Z \geq 0, \qquad 0 < r \leq 1.$$

Here a, b are inherited mobile capacity moments, D_L is the coefficient determined by the common available labor and production primitives, Z is the coerced-supply coefficient, and r is the certified producer's retained fraction after the existing uniform charge. Zero Z is the no-coerced-supply boundary of these market equations, not an assertion that the paper's application has no coerced mines.

Let p_0 be the pooled price, $z = p_C/p_D > 0$, and $d = p_D/p_0$. Define

$$X = a(rz)^{mk}, \qquad q_C(z) = \frac{D_L X}{rz(X+b)^{\alpha}}, \qquad q_D(z) = \frac{D_L b}{(X+b)^{\alpha}} + Z. \tag{144}$$

Actual quantities are $Q_j = p_D^{m-1} q_j(z)$. Hence

$$\begin{aligned} F(z) &= zq_C(z) + q_D(z) = D_L \frac{X/r + b}{(X+b)^{\alpha}} + Z, \\ F_0 &= D_L(a+b)^{1-\alpha} + Z, \qquad p_D^m = \frac{I}{F(z)}, \qquad p_0^m = \frac{I}{F_0}, \qquad d^m = \frac{F_0}{F(z)}. \end{aligned} \tag{145}$$

The pooled quantities use $z = r = 1$, so $X = a$.

For the existing North–South CES demand, put $0 \le \nu < 1$, $\omega > 0$ and $\varepsilon > 1$, where ν is North's share of dated mineral spending. The market root satisfies

$$\mathcal{R}(z) := \frac{zq_C(z)}{q_D(z)} = \frac{D_L X}{r\{D_L b + Z(X+b)^{\alpha}\}} = \frac{\nu + \omega z^{1-\varepsilon}}{1-\nu}. \tag{146}$$

The root is unique and increases with ν. Indeed, X increases with z, and the derivative of the middle expression with respect to X is positive: its numerator, apart from positive factors, is

$$D_L b + Z(X+b)^{\alpha-1}\{b + (1-\alpha)X\} > 0.$$

The left side of (146) increases from zero to infinity, while the right side decreases from infinity to a finite limit. Increasing ν strictly raises the right side at a fixed z. The boundary $\nu = 1$ has zero dirty final demand and is not covered by division by $1 - \nu$.

It is useful to express a positive candidate root as a North-share threshold:

$$\mathcal{N}(z) = \frac{\mathcal{R}(z) - \omega z^{1-\varepsilon}}{1 + \mathcal{R}(z)}. \tag{147}$$

This map is strictly increasing. The equilibrium satisfies $\nu = \mathcal{N}(z)$; thus a crossing z_j gives an exact test $z > z_j$ if and only if $\nu > \mathcal{N}(z_j)$. A negative threshold simply means that every admissible nonnegative North share exceeds it. No negative North share is treated as an economic equilibrium.

K.1.2 A spending identity for the current ownership model

Let E_C^P denote pooled clean expenditure and E_C^C gross expenditure on certified output after segmentation. Define

$$\begin{aligned} C_0 &= F_0(a+b)^{\alpha} = D_L(a+b) + Z(a+b)^{\alpha}, \\ \mathcal{T}(X) &= D_L(X/r + b) + Z(X+b)^{\alpha}, \qquad s_0 = \frac{aD_L}{C_0} \in (0,1). \end{aligned} \tag{148}$$

The symbol s_0 here is the pooled clean *spending share*, not research effort.

Proposition K.1 (Exact spending and operating-flow identities)**.** *In the matched market above, the certified inventor's first licensed machine-profit flow and the mobile dirty owner's proportional assigned prototype gain have respective multipliers*

$$M_C = \frac{XC_0}{a\mathcal{T}(X)}, \qquad M_D = \frac{C_0}{\mathcal{T}(X)}. \tag{149}$$

In particular,

$$\frac{E_C^P}{I} = s_0, \qquad M_C = \frac{rE_C^C}{E_C^P}, \qquad \frac{M_C}{M_D} = \frac{X}{a}. \tag{150}$$

The dirty multiplier requires the current model's proportional owned prototype claim, unchanged claim share and homogeneous success and failure operating problems on the same network. It is not a statement about arbitrary changes in ownership, entry or network membership.

Proof. The current labor-clearing equation gives the ordinary-wage ratio

$$W = \frac{w_C^L}{w_P^L} = d^m \left(\frac{X+b}{a+b}\right)^{1/k}.$$

The existing matched-owner calculation gives $M_D = d^{mk} W^{-\alpha k}$ and $M_C = (rz)^{mk} M_D$. Since $k(1-\alpha) = 1$,

$$M_D = d^m \left(\frac{a+b}{X+b}\right)^{\alpha} = \frac{F_0(a+b)^\alpha}{F(z)(X+b)^\alpha} = \frac{C_0}{\mathcal{T}(X)}.$$

Using $(rz)^{mk} = X/a$ gives the first identity in (149). Direct expenditure accounting gives

$$E_C^P = p_0^m \frac{D_L a}{(a+b)^\alpha}, \qquad rE_C^C = p_D^m \frac{D_L X}{(X+b)^\alpha}.$$

Dividing the first expression by $I = p_0^m F_0$ proves $E_C^P/I = s_0$. Dividing the second by the first proves the net-spending identity. □

Proposition K.2 (Sharp operating-flow crossings)**.** *As functions of $X > 0$, M_D is strictly decreasing and M_C is strictly increasing. There is a unique $X_D > 0$ such that $M_D = 1$, characterized by*

$$\mathcal{T}(X_D) = C_0, \qquad X_D \le a. \tag{151}$$

A finite crossing $M_C = 1$, and hence a finite state with $M_C > 1$, exists if and only if

$$r > s_0. \tag{152}$$

When it exists it is unique, satisfies $X_C \ge a$, and solves

$$X_C C_0 = a\mathcal{T}(X_C). \tag{153}$$

For $r < 1$ the inequalities $X_D < a < X_C$ are strict. For $r = 1$ both crossings equal a. If $r \le s_0$, $M_C < 1$ at every finite X; at $r = s_0$ it tends to one from below.

Proof. Both $\mathcal{T}$ and $\mathcal{T}'$ are positive. Thus M_D strictly decreases. Furthermore,

$$\mathcal{T}(X) - X\mathcal{T}'(X) = D_L b + Z(X+b)^{\alpha-1}\{b + (1-\alpha)X\} > 0,$$

so differentiating $XC_0/[a\mathcal{T}(X)]$ proves strict increase of M_C. At zero, $\mathcal{T}(0) = D_L b + Zb^\alpha < C_0$, and $\mathcal{T}(X) \to \infty$. This proves existence and uniqueness of X_D. At $X = a$,

$$\mathcal{T}(a) = C_0 + aD_L(1/r - 1), \qquad M_C(a) = M_D(a) = \frac{C_0}{C_0 + aD_L(1/r-1)} \le 1. \tag{154}$$

This locates X_D and any certified crossing. Finally,

$$\lim_{X \downarrow 0} M_C(X) = 0, \qquad \lim_{X\to\infty} M_C(X) = \frac{rC_0}{aD_L} = \frac{r}{s_0},$$

where $\alpha < 1$ makes $Z(X+b)^\alpha = o(X)$. Continuity and strict increase give all remaining statements. If $r > s_0$, sufficiently large z is an admissible market root for some $\nu < 1$, since $\mathcal{N}(z) \to 1$ from below. The possibility of strict improvement is therefore not an artifact of allowing negative North shares. The equality crossing itself may correspond to a negative threshold, in which case it is outside the admissible North-share domain and the strict improvement already holds throughout that domain. □

Corollary K.3 (What the net-spending condition preserves)**.** *If $rE_C^C \ge E_C^P$, the certified licensed first operating flow weakly rises, the mobile dirty proportional operating gain weakly falls, and the coerced fixed-capacity proportional operating gain weakly falls. For a positive charge $r < 1$, both dirty multipliers are strictly below one. These are operating-flow comparisons; they do not by themselves compare complete research prizes, grant requirements or two dynamic equilibria.*

Proof. By Proposition K.1, the hypothesis is $M_C \ge 1$. Proposition K.2 implies $X \ge a$ and therefore $M_D \le 1$, with strict inequalities $X > a$ and $M_D < 1$ when $r < 1$. The function F in (145), viewed as a function of X, is strictly increasing because

$$\frac{d}{dX}\{(X/r + b)(X+b)^{-\alpha}\} = (X+b)^{-\alpha-1}\{(1-\alpha)X/r + b(1/r - \alpha)\} > 0.$$

Also $F(a) \ge F_0$, strictly if $r < 1$. Hence $d^m = F_0/F \le 1$. The current fixed-capacity coerced branch has operating-profit and proportional prototype-gain multiplier d^m, which proves its claim. Multiplying by a zero baseline flow gives equality even when the relevant multiplier is strictly below one. □

Remark K.4 (The complete-value comparison remains necessary). The identity $M_C = rE_C^C/E_C^P$ recovers an economically interpretable spending condition for current-model first flows without converting the certified patent into a one-period asset. A larger funded

installation bill, a different surviving licensed opportunity, additional verification use or a laboratory-cost change can nevertheless reverse the complete research comparison. Those quantities remain in the exact support threshold; this corollary does not set them to zero or assume them equal.

K.1.3 Price and operating thresholds are distinct

There is a unique positive $\bar{X}$ at which $F = F_0$, and write $\bar{z} = r^{-1}(\bar{X}/a)^{1/(mk)}$. The dirty-price test is

$$p_D < p_0 \quad \Longleftrightarrow \quad X > \bar{X} \quad \Longleftrightarrow \quad \nu > \bar{\nu} := \mathcal{N}(\bar{z}). \tag{155}$$

This is the current mixed-site test. It does not use the earlier independent-power price threshold.

Proposition K.5 (Location of the price–prize reversal interval)**.** *For $r < 1$,*

$$ar^{mk} < \bar{X} < X_D < a, \qquad 1 < \bar{z} < 1/r. \tag{156}$$

Consequently, at any admissible root with $\bar{X} < X < X_D$, the dirty price falls while the mobile dirty proportional operating gain rises. If $r > s_0$, the additional certified-flow crossing satisfies $X_C > a$. For $r = 1$, the price and both operating-flow crossings all equal $X = a$.

Proof. The derivative displayed in the preceding proof establishes strict increase of F in X. Its lower endpoint is $D_L b^{1-\alpha} + Z < F_0$ and its upper endpoint is infinity. At $z = 1$, put $n = mk$ and

$$g(r) = \frac{ar^{n-1} + b}{(ar^n + b)^{\alpha}}.$$

Its derivative has the sign of

$$(m-1)ar^n + b\{n - 1 - \alpha nr\} > 0,$$

because $n - 1 - \alpha nr \geq m - 1 > 0$ for $r \leq 1$. Thus for $r < 1$, $F(1) = D_L g(r) + Z < D_L g(1) + Z = F_0$. At $z = 1/r$, $X = a$ and $F > F_0$. This proves the location of $\bar{X}$ and $\bar{z}$. At the price crossing $d = 1$, so

$$M_D(\bar{X}) = \left(\frac{a+b}{\bar{X}+b}\right)^{\alpha} > 1.$$

The strictly decreasing M_D therefore reaches one only at $X_D > \bar{X}$. The remainder follows from Proposition K.2. When $r = 1$, substitution gives all three crossings at $X = a$. □

The interval in Proposition K.5 describes the current model's reversal precisely; it is not a numerical anomaly. A threshold may lie outside the economically admissible North-share range. Its implication is applied only to roots whose $\mathcal{N}(z)$ belongs to $[0, 1)$.

K.1.4 Closed forms within the current specification

The following formulas either hold throughout the stated current snapshot or evaluate parameter cases of that same specification. Choosing a parameter case is not evidence that the general mixed-site equilibrium has an elementary closed form.

Corollary K.6 (Zero-charge North-share threshold)**.** *For $r = 1$, the common dirty-price and operating-flow crossing is $z = 1$, and its exact North-share threshold is*

$$\bar{\nu} = (1+\omega)s_0 - \omega = \frac{D_L a - \omega\{D_L b + Z(a+b)^{\alpha}\}}{D_L(a+b) + Z(a+b)^{\alpha}}. \tag{157}$$

Thus $\nu > \bar{\nu}$ simultaneously gives $p_D < p_0$, $M_C > 1$ and $M_D < 1$. Equality gives equality in all three comparisons, whenever that North share is admissible.

Proof. At $r = z = 1$, gross clean expenditure has share s_0, so $\mathcal{R}(1) = s_0/(1-s_0)$. Substitution into (147) gives (157). Monotonicity and the coincident crossings establish the signs. □

Proposition K.7 (Positive-charge price crossing when $\alpha = 1/2$)**.** *In the existing parameter case $\alpha = 1/2$ (hence $k = 2$), including positive Z, the mixed-site price crossing is*

$$\begin{aligned} \bar{y} &= \frac{r\sqrt{a+b} + \sqrt{r^2(a+b) + 4b(1-r)}}{2}, \\ \bar{X} &= \bar{y}^2 - b, \qquad \bar{z} = \frac{1}{r}\left(\frac{\bar{X}}{a}\right)^{1/(2m)}, \\ \bar{\mathcal{R}} &= \frac{D_L \bar{X}}{r(D_L b + Z\bar{y})}, \qquad \bar{\nu} = \frac{\bar{\mathcal{R}} - \omega \bar{z}^{1-\varepsilon}}{1+\bar{\mathcal{R}}}. \end{aligned} \tag{158}$$

For $0 < r < 1$, $\bar{X} > 0$ and $1 < \bar{z} < 1/r$.

Proof. The equation $F = F_0$ cancels Z and D_L, leaving $(X/r + b)/\sqrt{X+b} = \sqrt{a+b}$. Substituting $y = \sqrt{X+b}$ yields

$$y^2 - r\sqrt{a+b}\, y - b(1-r) = 0.$$

The positive root is the displayed $\bar{y}$. Existence and location already proved for the unique crossing ensure $\bar{y} > \sqrt{b}$ and $\bar{X} > 0$; the transformations to $\bar{z}$ and $\bar{\nu}$ then follow from (144) and (147). At $r = 1$ this formula reduces to $\bar{X} = a$, $\bar{z} = 1$. □

Proposition K.8 (Explicit spending crossing when $\alpha = 1/2$)**.** *For $\alpha = 1/2$ and $r > s_0$, let $A = C_0 - aD_L/r > 0$. The unique crossing of certified producers' net expenditure with pooled clean expenditure is*

$$\begin{aligned} y_C &= \frac{aZ + \sqrt{a^2Z^2 + 4Ab(A + aD_L)}}{2A}, \\ X_C &= y_C^2 - b, \qquad z_C = \frac{1}{r}\left(\frac{X_C}{a}\right)^{1/(2m)}, \qquad \nu_C = \mathcal{N}(z_C). \end{aligned} \tag{159}$$

Thus $rE_C^C > E_C^P$ if and only if $\nu > \nu_C$ at an admissible market root. The weak inequality uses $\nu \geq \nu_C$.

Proof. The equation $XC_0 = a\mathcal{T}(X)$ is $AX = aD_Lb + aZ\sqrt{X+b}$. With $y = \sqrt{X+b}$ it becomes

$$Ay^2 - aZy - b(A + aD_L) = 0.$$

Its constant term is negative and its leading coefficient positive, so there is exactly one positive root, given in (159). The unique positive X crossing from Proposition K.2 identifies this root with y_C. Strict monotonicity of both M_C and the market-root map proves the test. □

Corollary K.9 (No-coerced-supply boundary)**.** *At the existing boundary $Z = 0$, for any $0 < \alpha < 1$, $s_0 = a/(a+b)$. If $r > s_0$, the certified spending crossing and mobile dirty-flow crossing have the elementary expressions*

$$X_C = \frac{abr}{r(a+b) - a}, \qquad X_D = ar, \qquad z_C = \frac{1}{r}\left(\frac{X_C}{a}\right)^{1/(mk)}. \tag{160}$$

The North-share test is again $\nu > \mathcal{N}(z_C)$. These formulas are boundary evaluations; positive coerced capacity must retain its Z term, as in Propositions K.7 and K.8.

Proof. Now $C_0 = D_L(a+b)$ and $\mathcal{T}(X) = D_L(X/r + b)$. Solving $XC_0 = a\mathcal{T}(X)$ gives the first expression, whose denominator is positive exactly when $r > s_0$. Solving $\mathcal{T}(X) = C_0$ gives $X_D = ar$. The transformation to z_C is its definition. □

Remark K.10 (Scope of the closed forms). The conditional multipliers and all threshold tests retain the actual current-model market root. For arbitrary α, the equations defining $\bar{X}$ and X_C are exact, monotone scalar characterizations; no general elementary radical formula is asserted. Unequal productive labor, endogenous changes in active networks, nonlinear selling continuations, or a changed funding or ownership specification require their actual market equations and cannot inherit these equal-labor formulas merely by reusing their notation.

K.2 Complete rewards and the preservation of private research

The spending identity concerns an operating flow. A private research decision uses the owner's complete conditional-success value and the actual marginal cost of scientific effort. This section gives an exact bridge between the two, without holding costs fixed by assumption.

K.2.1 Actual effort, values and support

Let $x \in \{P, C\}$ denote the pooled and certified regimes, and use a common positive date-t monetary normalization c. Keep the existing project masses, total scientist mass S, success technology

$$h_i(s) = 1 - e^{-\nu_i^R s}, \qquad \nu_i^R > 0,$$

and actual verification allocation $S_V^x < S$. Full-certified effort per project is $s_x = (S - S_V^x)/N_C > 0$. Ordinary labor used by verification is separate from these scientists; the equal ordinary-labor condition in Section K.1 does not imply equal s_x.

Write $\widehat{J}_i^x = J_i^x/c$ for complete date-t conditional-success values, excluding current scientist wages and laboratory costs. Let R_i^x be the actual normalized laboratory cost: it is differentiable, convex and nondecreasing, and $R_i^x(0) = 0$. In particular, it includes any ordinary-labor cost at the actual wage. The nonnegative conditional-success payment G_i^x is in the same normalized date-t units; its promised next-date cash value is cG_i^x/δ_R. Define

$$\begin{aligned} B_x &= \max\left\{0, \max_{i \neq C}\left(h_i'(0)[\widehat{J}_i^x + G_i^x]_+ - R_i^{x\prime}(0)\right)\right\}, \\ \Lambda_x &= \frac{R_C^{x\prime}(s_x) + B_x}{h_C'(s_x)} = \frac{e^{\nu_C^R s_x}}{\nu_C^R}\left(R_C^{x\prime}(s_x) + B_x\right), \\ \Gamma_x &= \Lambda_x - \widehat{J}_C^x, \qquad G_{C,x}^{\min} = [\Gamma_x]_+. \end{aligned} \tag{161}$$

Here $[y]_+ = \max\{y, 0\}$; inactivity is explicitly among the rival choices. Each set of rewards must be valued on a public continuation consistent with the allocation being assessed. Finite complete rewards, the existing scientist-market hypotheses, actual funding and admissible deployment remain required.

Lemma K.11 (Exact support criterion; existing result)**.** *A nonnegative additive certified success payment weakly supports the proposed full-certified allocation if and only if $G_C^x \geq [\Gamma_x]_+$. Strict marginal dominance over inactivity and every rival requires and is equivalent to*

$$G_C^x \geq 0, \qquad G_C^x > \Gamma_x.$$

Thus zero certified support gives strict private dominance exactly when $\Gamma_x < 0$. To call the entire research allocation private, set all rival research payments to zero as well.

Proof. At full-certified effort the scientist wage that supports an interior certified project is

$$w_x^R = h_C'(s_x)(\widehat{J}_C^x + G_C^x) - R_C^{x\prime}(s_x).$$

When the success value is nonnegative, diminishing h_C' and convex laboratory costs make the project's net objective concave; its first-order condition supports this effort. With zero success value and zero marginal costs the same condition gives a weakly optimal allocation. A negative success value with nonnegative costs and wage cannot support positive effort. The condition $w_x^R \geq B_x \geq 0$ automatically excludes that case, since $h_C'(s_x) > 0$.

For a rival with nonnegative success value, its marginal return at zero is its largest marginal return. If its success value is negative, every positive effort has negative success contribution and nonnegative costs, so zero effort is optimal; the positive part in B_x safely implements that case. Thus $w_x^R \geq B_x$ is exactly the weak condition for no rival or inactivity to improve the proposed allocation. Rearranging gives $G_C^x \geq \Gamma_x$. Intersecting with nonnegative payments proves the first statement. Replacing the marginal comparison by a strict one proves the second. This is a conditional research-market calculation, not a new existence theorem for the aggregate economy. □

K.2.2 Exact preservation and cost headroom

For any object with regime superscripts define $\Delta Y = Y^C - Y^P$. Use $\Delta\Lambda = \Lambda_C - \Lambda_P$, and set

$$\mathcal{D} = \Delta\widehat{J}_C - \Delta\Lambda.$$

The symbol $\mathcal{D}$ in this section denotes a scalar improvement in the support comparison, not the invariant state domain used later.

Theorem K.12 (Exact matched-state preservation)**.** *At the consistently valued matched states,*

$$\Gamma_C = \Gamma_P - \mathcal{D}, \qquad G_{C,C}^{\min} = [\Gamma_P - \mathcal{D}]_+. \tag{162}$$

Strict private full-certified research in the certified regime holds exactly when $\mathcal{D} > \Gamma_P$. If the pooled state is already strictly private, with margin $m_P = -\Gamma_P > 0$, it remains strictly private exactly when

$$\mathcal{D} > -m_P. \tag{163}$$

The stronger condition $\mathcal{D} \geq 0$ weakly lowers the untruncated threshold and preserves every strictly private pooled state to which that condition applies. When $\mathcal{D} \geq 0$, the weak-support grant falls by exactly

$$G_{C,P}^{\min} - G_{C,C}^{\min} = \min\{[\Gamma_P]_+, \mathcal{D}\}. \tag{164}$$

Proof. Subtract the two definitions of Γ_x in (161). Strict private dominance is $\Gamma_C < 0$, which gives the first inequality and, when $\Gamma_P = -m_P$, (163). The weak order of the untruncated thresholds is equivalent to $\mathcal{D} \geq 0$. Finally, for any real u and $v \geq 0$, $[u]_+ - [u - v]_+ = \min\{[u]_+, v\}$; check $u \leq 0$, $0 < u \leq v$, and $u > v$. Apply this identity with $u = \Gamma_P$, $v = \mathcal{D}$. □

Let π^P be the pooled first licensed operating flow in the matched production-date comparison. Denote by ΔV_{tail} the actual date-t difference in subsequent survival-weighted operating values. It is zero only when those later opportunities are explicitly matched. For a stochastic comparison, use the corresponding conditional expected first-flow difference instead of factoring an endogenous random multiplier out of an expectation. Under the fixed snapshot multipliers of Section K.1, put

$$\Delta O_C = \delta_R(M_C - 1)\pi^P + \Delta V_{\text{tail}}. \tag{165}$$

If $\mathcal{C}_C^x$ is the full discounted incremental contractual cost, including refunds with their actual signs, then

$$\Delta J_C = \Delta O_C - \Delta \mathcal{C}_C.$$

Installation, access, collateral, settlements and residual claims are all included. A fee already deducted from the production receipt is not deducted again.

Corollary K.13 (Exact additional-cost allowance)**.** *For a strictly private pooled state of margin m_P, strict preservation is equivalent to*

$$\boxed{\Delta \mathcal{C}_C < \Delta O_C - c\,\Delta\Lambda + c\,m_P.} \tag{166}$$

At equality the certified comparison is a tie. The stronger condition for weakly lowering the untruncated support threshold is exactly

$$\Delta \mathcal{C}_C \leq \Delta O_C - c\,\Delta\Lambda. \tag{167}$$

These conditions retain actual changes in verification, available research effort, laboratory costs and competing returns through $\Delta\Lambda$.

Proof. Substitute $\Delta \widehat{J}_C = (\Delta O_C - \Delta \mathcal{C}_C)/c$ into (163) and multiply by $c > 0$. Equality makes $\mathcal{D} = -m_P = \Gamma_P$, hence $\Gamma_C = 0$. Repeating the rearrangement for $\mathcal{D} \geq 0$ proves (167). □

The net-spending condition established in Section K.1 ensures that the first term of (165) is nonnegative and that the relevant dirty operating gains do not rise. It supplies an economic channel for satisfying (166); it does not determine its remaining terms. An increase in ordinary wages can raise laboratory costs, a verification allocation can change s_x, and changing research incentives can change later patent replacement and licensing. All of these enter the exact test. Even $\Gamma_C < \Gamma_P$ cannot strictly reduce a grant that is already zero.

K.2.3 The existing reversal examples

For reproducibility retain the manuscript's admitted snapshot

$$m = 2, \quad \alpha = \tfrac{1}{2}, \quad k = 2, \quad a = b = D_L = Z = 1, \quad I = \sqrt{2} + 1, \quad \varepsilon = 2.$$

It normalizes $p_0 = 1$. For each displayed r, z, ω, use the market-clearing equations of Section K.1; the Northern share is chosen as

$$\nu = \frac{z q_C(z)/q_D(z) - \omega z^{1-\varepsilon}}{1 + z q_C(z)/q_D(z)}.$$

Both shares lie in $[0, 1)$. The table displays rounded diagnostics, not calibrated economies.

Case	r	z	ω	ν	d	W
Favorable flows	1	2	1	.636197	.686469	1.373890
Price–gain reversal	.8	1.14	.1	.209428	.995997	.912379

Case	M_C	Mobile M_D	Coerced multiplier d^2
Favorable flows	2.586145	.161634	.471240
Price–gain reversal	.746167	1.078591	.992010

In the second row dirty price falls, but the mobile gain rises because ordinary labor becomes cheaper. The coerced gain falls instead. The certified net-spending condition fails in that row, so there is no conflict with the spending theorem.

The first row also supports the manuscript's complete-cost reversal. Set $c = 1$, $\delta_R = .9$, retain the same later operating opportunities, and write $p_* = \pi^P > 0$. Choose $j > 0$ below the pooled gross operating value and nonnegative baseline contractual costs giving $J_C^P = j$. The existing admissible additional nonrefundable deployment bill has date-t value

$$\Delta\mathcal{C}_C = .9(M_C - 1)p_* + 2j.$$

Thus $\Delta J_C = -2j$, so $J_C^C = -j$. With zero laboratory marginal costs, no rival payments and actual dirty costs making both dirty rewards nonpositive, $B_P = B_C = 0$, giving

$$\Gamma_P = -j, \qquad \Gamma_C = j.$$

The weak-support grant rises from zero to j, notwithstanding the favorable certified operating flow. A payment strictly above j is needed for strict dominance in the second regime. This is the existing finite-bill contract example: funded verification, a positive production-labor remainder, admissible machine quality, actual participation and outside resources covering the stated bills remain part of its feasibility conditions. It is not a claim about the empirical cost of certification or a constructed pair of full dynamic equilibria.

K.3 Sharper path dependence within the existing model

This section tightens the current model's sufficient private-takeover certificate. It does not replace surviving certified patents by one-period claims, hold world expenditure fixed along growth, or introduce a new production or research technology. The improvement comes from using the existing bounds jointly, rather than bounding three error terms separately. The resulting cutoff remains sufficient for a selected continuation; it is not a necessary economic threshold below which an economy is trapped.

K.3.1 Existing hypotheses and the exact marginal comparison

Use the homogeneous two-type application already stated in the current path-dependence theorem. There is one unit of certified projects and one unit of dirty projects, a fixed scientist mass S, and a funded verification provision $S_V < S$. Write

$$\begin{aligned} s^* &= S - S_V > 0, \qquad h_i(s) = 1 - e^{-\nu_i^R s}, \\ R_i(s) &= \frac{r_i s^2}{2}, \qquad r_i \geq 0, \quad \nu_i^R > 0, \\ H_C &= h_C'(s^*) = \nu_C^R e^{-\nu_C^R s^*} > 0. \end{aligned} \tag{168}$$

Here r_C is the laboratory-cost coefficient, distinct from the retained receipt fraction r in Section K.1; $\mathcal{A}_C$ below is a marginal research return, distinct from the earlier flow multiplier M_C. The certification-directed allocation has $s_C = s^*$ and $s_D = 0$. The normalization is the inherited certified technology $A = A_{C,t} > 0$: $\widehat{J}_i = J_i/A$. These are complete conditional-success owner values, including the actual dated contractual deductions. A research success probability is not multiplied into J_i a second time.

The comparison here sets the dirty research payment to zero. It concerns private takeover of the full-certified allocation, not a situation in which the competing project continues to receive a public research payment. At the actual consistent values, define

$$\mathcal{A}_C = H_C \widehat{J}_C - r_C s^*, \qquad \mathcal{B} = \nu_D^R [\widehat{J}_D]_+, \qquad \Gamma = \frac{r_C s^* + \mathcal{B}}{H_C} - \widehat{J}_C, \qquad G_C^{\min} = [\Gamma]_+. \tag{169}$$

Here $[x]_+ = \max\{x, 0\}$, and inactivity is included because $\mathcal{B} \geq 0$. Quadratic dirty laboratory costs give $R_D'(0) = 0$. The current model's exact support condition is $G_C \geq G_C^{\min}$ for weak support, and $G_C \geq 0$, $G_C > \Gamma$ for strict marginal dominance. In particular, strict private takeover is exactly $\mathcal{A}_C > \mathcal{B}$, or $\Gamma < 0$, at these consistent values. This research-market statement does not by itself supply a joint market, government and research equilibrium.

Retain the current theorem's forward-invariant domain $\mathcal{D}$ of technologies, active licensed networks, adoption states, outstanding claims, verification workloads and funding states. On this domain, for $A \geq A_{\min} > 0$, its complete-reward bounds are

$$\widehat{J}_C \geq j_0 - \frac{K_C}{A}, \qquad [\widehat{J}_D]_+ \leq \frac{D_0}{A} + \frac{D_1}{A^q}, \qquad q = 1 - \theta_\Sigma \in (0,1), \qquad K_C, D_0, D_1 \geq 0, \tag{170}$$

where

$$j_0 = \delta_R \beta_m (1 + \gamma_C) \underline{\mathcal{M}} - \kappa_C, \qquad m_\infty = H_C j_0 - r_C s^* > 0. \tag{171}$$

The growth application has $\theta_\Sigma = [1 + \beta_m(1 - \alpha_\Sigma)(\varepsilon_M - 1)]^{-1}$, with $0 < \alpha_\Sigma < 1$ and $\varepsilon_M > 1$; its final-producer closure is distinct from the dated fixed-budget comparison. The discount factor is $0 < \delta_R < 1$; β_m is the actual admissible unit-quote machine margin in the stated production application; $\gamma_C > 0$ is the certified quality step; $\underline{\mathcal{M}} > 0$ is the existing licensed-market multiplier bound; and $\mathcal{C}_C \leq \kappa_C A + K_C$ is the existing bound on contractual deductions. If the unit-quote branch is inadmissible, the original paper

requires the actual constrained machine margin: the displayed formula must not be used with an inapplicable β_m. The dirty bound includes the actual one-period gain and all remaining net claims, including the finite diffusion pipeline. It is not a bound on gross sales alone.

The underlying domain and its invariance are substantive hypotheses already required by the current theorem. The bounds must apply to the actual unsubsidized future values after withdrawal. A favorable calculation using values that still anticipate future research payments is not a withdrawal certificate.

K.3.2 A tighter sufficient cutoff and an explicit grant bound

Proposition K.14 (Combined reward-envelope certificate)**.** *Under* (168)–(171)*, put*

$$\mathfrak{a} = H_C K_C + \nu_D^R D_0, \qquad \mathfrak{b} = \nu_D^R D_1, \qquad f(A) = \frac{\mathfrak{a}}{A} + \frac{\mathfrak{b}}{A^q}. \tag{172}$$

Then, throughout the stated domain,

$$\mathcal{A}_C - \mathcal{B} \geq m_\infty - f(A), \qquad \Gamma \leq \frac{f(A) - m_\infty}{H_C}, \qquad 0 \leq G_C^{\min} \leq \frac{[f(A) - m_\infty]_+}{H_C}. \tag{173}$$

Consequently $f(A) < m_\infty$ is sufficient for strict private takeover. It is the exact strict test furnished by these combined bounds; it is not a necessary condition on the actual private rewards.

Proof. The certified envelope gives $\mathcal{A}_C \geq m_\infty - H_C K_C / A$. The dirty envelope and $R_D'(0) = 0$ give $\mathcal{B} \leq \nu_D^R (D_0/A + D_1/A^q)$. Subtracting proves the first inequality. Since $H_C \Gamma = \mathcal{B} - \mathcal{A}_C$ and $H_C > 0$, the second follows. Applying the increasing positive-part function and using $[x/H_C]_+ = [x]_+ / H_C$ proves the third. Strict positivity of the lower bound in the first inequality implies the exact strict condition in (169). Failure of that lower-bound test says nothing about the sign of the actual marginal difference. □

For any chosen reporting margin $\eta \in (0, 1)$, define $c_\eta = (1 - \eta) m_\infty > 0$. If $\mathfrak{a} + \mathfrak{b} > 0$, let ϱ_η be the unique positive solution of

$$\frac{\mathfrak{a}}{\varrho_\eta} + \frac{\mathfrak{b}}{\varrho_\eta^q} = c_\eta. \tag{174}$$

If $\mathfrak{a} = \mathfrak{b} = 0$, set $\varrho_\eta = 0$ by convention; no positive root is claimed in this degenerate case. Define

$$A_\eta = \max\{A_{\min}, \varrho_\eta\}. \tag{175}$$

The freely chosen η specifies how much of an already proved strict margin is retained in the certificate; it is not an economic primitive.

Theorem K.15 (Refined invariant private region)**.** *For every maintained-domain state with $A \geq A_\eta$,*

$$\mathcal{A}_C - \mathcal{B} \geq \eta m_\infty > 0. \tag{176}$$

If the selected unsubsidized full-certified continuation stays in $\mathcal{D}$ and has nondecreasing A, the set of such states with $A \geq A_\eta$ is forward invariant for that continuation.

In particular, take $\eta = 1/4$. The refined cutoff satisfies

$$A_{1/4} \leq A^*_{\text{old}}, \qquad A^*_{\text{old}} = \max\left\{A_{\min}, \frac{4H_C K_C}{m_\infty}, \frac{4\nu_D^R D_0}{m_\infty}, \left(\frac{4\nu_D^R D_1}{m_\infty}\right)^{1/q}\right\}. \tag{177}$$

Thus the refinement weakly lowers the published sufficient technology cutoff while retaining its positive margin $m_\infty/4$.

Proof. If at least one coefficient is positive, f is continuous and strictly decreasing on $(0, \infty)$, tends to infinity at zero, and tends to zero at infinity. These facts prove existence and uniqueness in (174). For $A \geq A_\eta$, $f(A) \leq c_\eta$, so (173) gives (176). If both coefficients vanish, $f \equiv 0$ and the same conclusion follows directly. On the selected continuation, domain preservation and nondecreasing technology preserve $A \geq A_\eta$. The same strict comparison therefore holds at every next state. Induction proves forward invariance; it does not prove existence or uniqueness of that joint continuation.

At $A = A^*_{\text{old}}$, each of $H_C K_C/A$, $\nu_D^R D_0/A$, and $\nu_D^R D_1/A^q$ is at most $m_\infty/4$. Their sum is at most $3m_\infty/4 = c_{1/4}$. If the root is positive, monotonicity implies $\varrho_{1/4} \leq A^*_{\text{old}}$; the degenerate case is immediate. Since $A_{\min} \leq A^*_{\text{old}}$ as well, taking the maximum proves (177). Equality of the technology scale and the refined cutoff causes no tie, because the retained marginal margin is positive. □

Closed forms and what they do not assert. The following formulas solve the certificate, not the full market or the actual first date at which private research becomes viable:

$$\begin{aligned} \mathfrak{b} = 0: &\quad \varrho_\eta = \frac{\mathfrak{a}}{c_\eta}, \\ \mathfrak{a} = 0: &\quad \varrho_\eta = \left(\frac{\mathfrak{b}}{c_\eta}\right)^{1/q}, \\ q = \tfrac{1}{2}: &\quad \varrho_\eta = \left(\frac{\mathfrak{b} + \sqrt{\mathfrak{b}^2 + 4c_\eta \mathfrak{a}}}{2c_\eta}\right)^2. \end{aligned} \tag{178}$$

All formulas agree with the zero-root convention when both coefficients vanish. The first two follow by isolating the sole nonzero term. For $q = 1/2$, put $y = \sqrt{\varrho_\eta}$; multiplying (174) by y^2 gives $c_\eta y^2 - \mathfrak{b} y - \mathfrak{a} = 0$. Its positive root gives the displayed expression whenever a coefficient is nonzero; the zero case is treated by the stated convention. The half-exponent case is already permitted by the current model's parameter domain. It is an algebraic specialization, not an assumption imposed on the general economy. For arbitrary $q \in (0, 1)$, (174) remains a uniquely specified scalar root; no elementary closed form is asserted.

K.3.3 Withdrawal dates and stochastic hitting

Proposition K.16 (Refined deterministic withdrawal bound)**.** *Retain the current model's cohort-balanced assignment and success law under full-certified research, so that*

$$A_{C,t+1} = G_A A_{C,t}, \qquad G_A = 1 + \gamma_C h_C(s^*) > 1. \tag{179}$$

For $A_{C,0} \geq A_{\min}$, set

$$T_\eta = \max\left\{0, \left\lceil \frac{\log(A_\eta / A_{C,0})}{\log G_A} \right\rceil \right\}. \tag{180}$$

Along a consistent full-certified path satisfying the maintained market and financing conditions, funded payments strictly above the actual threshold before T_η can end at T_η. The resulting selected unsubsidized continuation is strictly viable. For $\eta = 1/4$, this guaranteed withdrawal date is no later than the published sufficient date.

Proof. Iteration of (179) gives $A_{C,t} = G_A^t A_{C,0}$. Because $\log G_A > 0$, (180) is the smallest nonnegative integer with $A_{C,t} \geq A_\eta$. At that date and thereafter, the strict margin and invariance theorem apply. Before that date, the exact support condition implements the stipulated allocation if its promised payments and the owners' separate expenses are funded. Monotonicity of the logarithm, ceiling, and nonnegative-part operations, together with (177), proves the comparison of dates. □

This construction evaluates complete patent values on the proposed continuation, including its actual unsubsidized future. The same research direction before and after withdrawal gives the same stated replacement hazard; it does not establish that a change in transfers leaves all prices, networks or budgets unchanged. Those objects must satisfy the original domain conditions along the proposed path. A payment $[f(A) - m_\infty]_+ / H_C$ is a sufficient *weak* incentive payment where the bounds apply; adding any strictly positive amount makes the comparison strict. Neither this bound nor a positive prospective patent value supplies the cash to pay laboratories, verification, deployment, or success-contingent transfers.

Proposition K.17 (Refined stochastic hitting bound)**.** *In the current finite stochastic application, condition on a fixed initial state with $A_{C,0} \geq A_{\min}$. Let $A_{C,t} > 0$ be adapted and nondecreasing, and let $\tau_\eta = \inf\{t \geq 0 : A_{C,t} \geq A_\eta\}$. Retain its existing condition that, conditional on every history before the target, a log-technology increment of at least $\ell > 0$ has probability at least $p \in (0, 1]$. Then*

$$\mathbb{P}(\tau_\eta < \infty) = 1, \qquad \mathbb{E}\tau_\eta \leq \frac{K_\eta}{p}, \qquad K_\eta = \left\lceil \frac{[\log(A_\eta / A_{C,0})]_+}{\ell} \right\rceil. \tag{181}$$

No independence of successive innovation events is needed.

Proof. If $K_\eta = 0$, the initial technology is already at least the target and $\tau_\eta = 0$. Otherwise augment the qualifying-event indicator at step t to equal one whenever the target has already been reached, and before hitting to equal $\mathbf{1}_{\{\log A_{C,t+1} - \log A_{C,t} \geq \ell\}}$. Its conditional

success probability is at least p at every step. Iterated conditional expectation therefore bounds the probability of n consecutive subsequent failures, conditional on any prior history, by $(1-p)^n$. The waiting time for the next augmented success is almost surely finite and has conditional expectation at most $\sum_{n\geq 0}(1-p)^n = 1/p$. Successive conditioning shows that the time of the K_ηth augmented success has expectation at most K_η/p and is finite almost surely. By that time either hitting occurred earlier, or K_η genuine increments of at least ℓ have occurred. Nondecreasing technology then guarantees that the target has been reached. Thus τ_η is bounded by that augmented-success time, proving both conclusions. □

For $p < 1$, this argument supplies no finite deterministic deadline: independent Bernoulli qualifying events, an allowed instance of the conditional probability condition, can fail for arbitrarily many successive dates with positive probability. Likewise finite expected hitting time does not finance every possible sequence of failed rounds. All original funding requirements continue until the actual stopping time.

K.3.4 The exact limiting trichotomy and the boundary case

For this subsection, the existing limiting-reward hypotheses replace the positive lower-margin certificate hypothesis $m_\infty > 0$. The negative and zero limiting cases are not asserted to lie in that positive-certificate domain: if its lower envelope also held with $m_\infty > 0$ along a path with $A_{C,t} \to \infty$, convergence would imply $j_\infty \geq j_0$ and hence $m_{\lim} \geq m_\infty > 0$.

Proposition K.18 (Limiting support classification)**.** *On a candidate full-certified path with fixed s^* and costs as above, suppose the complete normalized certified reward converges to a finite j_∞ and the competing marginal return satisfies $\mathcal{B}_t \to 0$. These are the limits used in the current model's continuing-support discussion, not consequences of technology growth alone. Define $m_{\lim} = H_C j_\infty - r_C s^*$. Then*

$$\Gamma_t \longrightarrow -\frac{m_{\lim}}{H_C}, \qquad G_{C,t}^{\min} \longrightarrow \left[-\frac{m_{\lim}}{H_C}\right]_+. \tag{182}$$

If $m_{\lim} > 0$, the actual marginal comparison is strictly private from some finite date onward. If $m_{\lim} < 0$, required support is eventually positive and has a strictly positive limit. If $m_{\lim} = 0$, the required payment tends to zero, but this does not determine whether strict private takeover occurs at a finite date.

Proof. Take limits in the exact identity $\Gamma_t = (r_C s^* + \mathcal{B}_t)/H_C - \widehat{J}_{C,t}$. Continuity of the positive-part function proves the second limit. A sequence with strictly negative limit is eventually negative; a sequence with strictly positive limit is eventually positive. These observations and (169) prove the first two cases. At zero limit, the exact finite-state identity is

$$\Gamma_t = \frac{\mathcal{B}_t}{H_C} + (j_\infty - \widehat{J}_{C,t}), \tag{183}$$

whose sign is not determined by its convergence to zero. In the existing manuscript's example $\widehat{J}_{C,t} = j_\infty - K/A_{C,t}$, $K > 0$, one has $G^{\min}_{C,t} \geq K/A_{C,t} > 0$ at every finite state, although this lower bound vanishes as technology grows. This establishes the asserted failure of a finite-withdrawal inference at the boundary. □

Here $m_{\lim}$ uses an actual limit, whereas m_∞ in (171) is computed from a uniform lower envelope. They must not be identified without a proof of equality. A nonpositive envelope coefficient does not imply a nonpositive actual limiting margin. Similarly, an initially favorable exact comparison is not a permanent withdrawal test unless the subsequent unsubsidized comparisons and domain conditions persist. The trichotomy does not assert a unique joint equilibrium, monotonicity of every actual reward, or a scalar necessary cutoff for states with different networks, pending innovations and claims.

K.3.5 Using surviving rents without changing patent ownership

The original reward bound uses only the first production flow of a successful certified vintage. Where the current model's maintained original licensed-core result applies to that vintage throughout survival, its later nonnegative flows supply a stronger bound. The next result uses that existing branch and its conditional quality-independent assignment. It cannot be applied merely because aggregate certified output stays positive, or because a new patent has a positive first-period market.

Proposition K.19 (Surviving-vintage reward bound)**.** *Maintain the fixed full-certified allocation from* (168)*, with one unit of certified projects and the manuscript's exclusive, quality-independent line assignment. Its endogenous replacement probability is*

$$\rho_C = h_C(s^*) \in (0, 1). \tag{184}$$

For the successful vintage v created at t, let a_v be its own fixed quality. Conditional on success and the date-t public state, $\mathbb{E}_t a_v = (1 + \gamma_C)A_{C,t}$. Let I_n indicate survival into $t + 1 + n$, with $I_0 = 1$. On the actual maintained-core continuation assume the bounds supplied by that existing result hold for this vintage while it survives:

$$0 < \underline{\mathcal{M}} \leq \mathcal{M}^v_{t+1+n} \leq \overline{\mathcal{M}} < \infty, \qquad \pi^v_{t+1+n} = \beta_m a_v \mathcal{M}^v_{t+1+n}. \tag{185}$$

Retain the original complete-cost bound $\mathcal{C}_{C,t} \leq \kappa_C A_{C,t} + K_C$ and finite actual contractual values. Then

$$\widehat{J}_{C,t} \geq j_0^{\text{surv}} - \frac{K_C}{A_{C,t}}, \qquad j_0^{\text{surv}} = \frac{\delta_R \beta_m (1 + \gamma_C)\underline{\mathcal{M}}}{1 - \delta_R(1 - \rho_C)} - \kappa_C. \tag{186}$$

No independence between the licensed multiplier and survival is needed.

Proof. The fixed allocation and the conditional quality-independent lottery give every surviving line replacement probability ρ_C at the next step, conditional on the current history, including its quality. Consequently, if $\mathcal{F}_{t+1+n}$ records that history,

$$\mathbb{E}[I_{n+1} \mid \mathcal{F}_{t+1+n}] = (1-\rho_C)I_n.$$

Since a_v is already known at that history, the tower property gives

$$\mathbb{E}_t[a_v I_{n+1}] = (1-\rho_C)\mathbb{E}_t[a_v I_n], \qquad \mathbb{E}_t[a_v I_n] = (1+\gamma_C)A_{C,t}(1-\rho_C)^n.$$

These are conditional-hazard calculations, not an independence assumption about common shocks, survival and market profits. On surviving histories, (185) bounds the multiplier pathwise. All operating flows are nonnegative, so summing their expectations is justified by monotone convergence. Their discounted sum is bounded below by

$$\delta_R\beta_m\underline{\mathcal{M}}(1+\gamma_C)A_{C,t}\sum_{n\geq 0}[\delta_R(1-\rho_C)]^n.$$

The upper multiplier bound gives the corresponding finite geometric upper bound; the ratio is less than one because $0 < \delta_R < 1$ and $0 \leq 1-\rho_C < 1$. Summing the lower bound and subtracting the actual contractual cost proves (186). No growth factor for the technological frontier enters this series: the old patent keeps its own quality a_v. □

The strengthened coefficient satisfies the explicit comparison

$$j_0^{\mathrm{surv}} - j_0 = \delta_R\beta_m(1+\gamma_C)\underline{\mathcal{M}}\frac{\delta_R(1-\rho_C)}{1-\delta_R(1-\rho_C)} \geq 0. \tag{187}$$

Accordingly the previous results also hold with $m_\infty^{\mathrm{surv}} = H_C j_0^{\mathrm{surv}} - r_C s^*$ whenever this quantity is positive and the surviving-core conditions hold on the same domain. The error coefficients $\mathfrak{a}, \mathfrak{b}$ do not change. Where the original $m_\infty > 0$, increasing its value weakly lowers the root and the resulting technology and date bounds for a fixed η. This may also provide a positive certificate when the first-flow bound alone is inconclusive; it does not make a disappearing licensed core admissible.

For reference, the manuscript already gives an exact constant-multiplier special case. If $\mathcal{M}_u^v = M$ throughout survival on the fixed allocation, its conditional operating value before contractual deductions is

$$\frac{\delta_R\beta_m(1+\gamma_C)A_{C,t}M}{1-\delta_R(1-\rho_C)}. \tag{188}$$

This is an exact special-case value, whereas (186) is a lower bound with a potentially varying multiplier. Neither formula makes a generic growing CES market stationary. The survival hazard remains generated by research effort; it is not replaced by an externally chosen displacement parameter.

Consistency before acceptance. All certificates in this section are conditional implications for the same model. Their economic application requires the original market and funding domain to support the proposed continuation, including the actual licensed old-vintage markets where the stronger bound is used. The proposed full-certified allocation generates the hazard used to value its patents; the resulting complete rewards must then satisfy the exact research condition at every date. This is a consistency check within an already supplied continuation, not an argument that assumes research is privately viable in order to prove its viability. Where a joint continuation or a maintained-core condition has not been established, the algebraic implication remains valid but its application remains conditional. No claim of a necessary lock-in cutoff, automatic funding, or a complete dynamic pooling counterfactual is made.

K.4 From the price shock to mine income and protection payments

K.4.1 An exact link for the coerced operation

This calculation uses the current model's already specified common-quote, installed-input coerced branch. Its curvature restriction is $\alpha_m = 1 - \beta_m$, where $m = 1/\beta_m > 1$; it is not imposed on the paper's other auction or production branches. At an uncontrolled site ℓ, fix positive installed capacity A_ℓ and procurement coefficient κ_ℓ, and write

$$z_\ell = A_\ell \kappa_\ell^{-(m-1)} > 0.$$

The current supply and optimized operating-surplus formulas are

$$Q_\ell(p) = z_\ell p^{m-1}, \qquad \Pi_\ell(p) = \frac{z_\ell p^m}{m}. \tag{189}$$

Procurement includes the stated household operating rent as well as its real input bill. The profit in (189) is therefore the operation's surplus after those procurement costs, not gross sales, the admission organizer's income, or a government payment.

Proposition K.20 (Exact transmission to the coerced mine)**.** *At the existing mixed-site snapshot, with capacity and procurement coefficient unchanged,*

$$\frac{\Pi_\ell^C}{\Pi_\ell^P} = \frac{F_0}{F(z)}, \qquad \frac{Q_\ell^C}{Q_\ell^P} = \left(\frac{F_0}{F(z)}\right)^{(m-1)/m}. \tag{190}$$

For a proportional, one-interval owned prototype improvement of size $\gamma_D > 0$, its gross incremental operating claim has the same multiplier $F_0/F(z)$. Consequently, on the existing threshold domain,

$$\nu > \bar{\nu} \iff p_D < p_0 \iff \Pi_\ell^C < \Pi_\ell^P \iff Q_\ell^C < Q_\ell^P.$$

Equality and the reversed inequalities correspond in the same way. The prototype statement concerns its operating claim before additional contractual differences.

Proof. Dividing the two formulas in (189) between regimes gives the multipliers d^{m-1} and d^m, where $d = p_D/p_0 > 0$. Market clearing gives $d^m = F_0/F(z)$, proving (190). Multiplying installed capacity by $1 + \gamma_D$ raises operating profit by $\gamma_D \Pi_\ell(p)$; its ratio is therefore the same. Since $m > 1$, both positive-power functions are strictly increasing in d. Apply the existing mixed-site price threshold, also derived in Section K.1. □

This supplies a direct current-model link between segmentation and the coerced mine. It does not extend the result to a general nonlinear auction continuation, a changing local procurement coefficient, a mine that changes entry status, or the mobile dirty operator. The last has the additional wage response in Section K.1. Complete dirty research value remains

$$J_D = \delta_R \, \mathbb{E}[\Delta \mathcal{O}_D] - \mathcal{C}_D,$$

so later settlements and other incremental contract costs still matter.

K.4.2 What can be concluded about protection income

Let $b_t^{\text{cash}} \geq 0$ be the actual protection flow while the site is uncontrolled. The current specification does not supply an equation mapping Π_ℓ to b^{cash}. In particular, it does not impose a constant profit share. The government's waiting payoff uses the actual payment.

Proposition K.21 (The limit of the level comparison)**.** *A fall in coerced operating profit does not, from the maintained model alone, imply either a fall or a rise in its actual protection payment. Even a restriction that the payment be no greater than operating profit would not determine the sign.*

Proof. The accounting restriction $0 \leq b^{\text{cash}} \leq \Pi$, if imposed, permits $\Pi^P = 10$ and $\Pi^C = 8$ with each of the payment pairs

$$(b^P, b^C) = (1, 2), \quad (2, 1), \quad (1, 1).$$

These satisfy a tighter restriction than a bounded multiple of dirty sales, yet display all three signs. They are accounting examples establishing nonidentification by the displayed restrictions, not alternative solved general equilibria. A particular sign requires a rule or further information linking the actual payments across states; that rule is absent from the maintained specification. □

Proposition K.22 (Relative decline under the existing growth bound)**.** *On a certified-directed growth path with $A_{C,t} \to \infty$, retain the current model's hypothesis that actual protection receipts are bounded by a fixed multiple of actual dirty sales plus bounded dated claims. Under its mixed CES growth conditions there are nonnegative constants B_0, B_1 such that*

$$0 \leq b_t^{\text{cash}} \leq B_0 + B_1 A_{C,t}^{\theta_\Sigma}, \qquad 0 < \theta_\Sigma < 1.$$

Consequently,

$$0 \leq \frac{b_t^{\text{cash}}}{A_{C,t}} \leq \frac{B_0}{A_{C,t}} + \frac{B_1}{A_{C,t}^{1-\theta_\Sigma}} \longrightarrow 0. \tag{191}$$

This does not establish a fall in the absolute payment. After successful control, the model sets the site's new local protection flow to zero from the next production date; earlier settlement claims remain.

Proof. The current growth result bounds coerced dirty sales by $O(A_C^{\theta_\Sigma})$ and mobile dirty sales by a strictly smaller exponent. A bounded multiple of those sales plus bounded dated claims has the stated envelope on the maintained domain, with finite constants as in the model's hypothesis. Dividing by $A_C > 0$ and using $\theta_\Sigma < 1$ proves (191). The upper bound allows an increasing sequence proportional to $A_C^{\theta_\Sigma}$, so it does not entail an absolute decrease. The cessation statement is the existing control transition: success changes local production at the next date; an attempt alone does not do so and does not erase historical claims. □

If an alternative monetary normalization B_t is uniformly comparable to $A_{C,t}$, the same relative-zero conclusion follows under that normalization. Without such a comparison it must not be asserted automatically. Nor does (191) alone imply intervention: the legal control value, feasible tax, cost and success of attempts, option to wait, convergence conditions and actual financing still enter the government's existing decision problem.

The exploratory conclusion is therefore precise. There is an exact static price-to-profit transmission on the coerced branch, an identified relative protection bound on the maintained growth path, and no identified static profit-to-payment sign. This appendix does not fill the missing sharing rule with an additional behavioral assumption.

K.5 Monotonicity beyond the installed-input specialization

Lemma K.23 (Selva output under a linear receipt)**.** *Fix the feasible Selva allocations and all household primitives. Suppose the armed group's continuation pays pQ for physical output Q, and the minimum implementation cost $C(Q)$, including subsistence and household incentive costs, is independent of p. Suppose optima exist. If $p_1 < p_2$, any optimal outputs Q_1 at p_1 and Q_2 at p_2 obey $Q_1 \leq Q_2$. No differentiability or uniqueness is needed.*

Proof. Optimality gives $p_1Q_1 - C(Q_1) \geq p_1Q_2 - C(Q_2)$ and $p_2Q_2 - C(Q_2) \geq p_2Q_1 - C(Q_1)$. Adding yields $(p_2 - p_1)(Q_2 - Q_1) \geq 0$; the first factor is positive. □

The argument permits a kinked implementation cost and endogenous household subsistence within that cost. It does not permit price-dependent outside options, a changing feasible set, or a nonlinear auction continuation to be silently substituted for the stipulated linear receipt. It signs optimal output, not the actual protection payment, whose level is not pinned down.

K.6 An exact intervention threshold with tax-dependent success

This extends the stationary calculation in the current reduced government problem. Normalize every dated flow by $B_t > 0$ and suppose $d = \delta_G B_{t+1}/B_t \in [0, 1)$ is constant. Let $b \geq 0$ be the normalized protection flow obtained by waiting, $k(\tau) \geq 0$ the current attempt cost, $P(\tau) \in (0, 1]$ the success probability, and $u(\tau)$ the normalized legal control value measured at the next control date. The feasible tax set $\mathcal{T}$ is nonempty and compact and k, P, u are continuous. It includes only actual private opening and funded implementation plans. Every continuation being compared, including repeated failure histories, is funded. Waiting and failure return to the same normalized uncontrolled state, with no binding balance distinction. These are conditions for the reduced Bellman equation, not consequences of positive expected value.

Proposition K.24 (Stationary intervention threshold)**.** *The reduced Bellman equation has the unique solution*

$$\begin{aligned} v &= \max\left\{b + dv, \max_{\tau\in\mathcal{T}}\big[-k(\tau) + dP(\tau)u(\tau) + d(1 - P(\tau))v\big]\right\} \\ &= \max\left\{\frac{b}{1-d}, \max_{\tau\in\mathcal{T}}\frac{dP(\tau)u(\tau) - k(\tau)}{1 - d + dP(\tau)}\right\}. \end{aligned} \tag{192}$$

An immediate attempt, repeated after failure at a maximizing tax, is strictly preferable to waiting exactly when

$$b < b^{\text{crit}}, \qquad b^{\text{crit}} = \max_{\tau\in\mathcal{T}}\frac{(1-d)[dP(\tau)u(\tau) - k(\tau)]}{1 - d + dP(\tau)}. \tag{193}$$

At equality, waiting and at least one maximizing attempt are optimal; other taxes need not tie. If $b^{\text{crit}} \leq 0$, no nonnegative protection flow makes intervention strictly preferable. If P is tax-independent, this reduces to

$$b^{\text{crit}} = \frac{1-d}{1-d+dP}\max_{\tau\in\mathcal{T}}\{dPu(\tau) - k(\tau)\}.$$

Proof. Every displayed action is an affine function $a_i + \lambda_i v$ with $0 \leq \lambda_i \leq d < 1$. Their maximum is a contraction with modulus at most d. Its unique fixed point is the largest action-specific fixed point $v^* = \max_i a_i/(1-\lambda_i)$: every action satisfies $a_i + \lambda_i v^* \leq v^*$, and an action attaining the largest fixed point gives equality. Compactness and continuity ensure that the tax maximum is attained. The waiting fixed point is $b/(1-d)$; the attempt fixed point is the fraction in (192). Comparing these fixed points gives (193). If the attempt value strictly exceeds the waiting fixed point, $b + dv^* < v^*$ and every optimal action is an attempt. The same state returns after failure, so repeating a maximizing attempt tax is optimal. Equality gives the stated indifference. Denominators are positive even when $d = 0$. □

Corollary K.25 (Paid catch-up in the intervention threshold)**.** *Suppose the funded control plan graduates after $H \geq 0$ periods. Let $r_a^G(\tau)$ be its expected normalized net public flow at control age a, and let $r_\infty^G(\tau)$ be its stationary mature flow. Then*

$$u(\tau) = \sum_{a=0}^{H-1} d^a r_a^G(\tau) + \frac{d^H r_\infty^G(\tau)}{1-d}, \tag{194}$$

and

$$b^{\text{crit}} = \max_{\tau \in \mathcal{T}} \frac{dP(\tau)\left[(1-d)\sum_{a=0}^{H-1} d^a r_a^G(\tau) + d^H r_\infty^G(\tau)\right] - (1-d)k(\tau)}{1-d+dP(\tau)}. \tag{195}$$

For a mature site with tax-independent P, k and optimized constant net public flow $r^{G,*}$, this becomes

$$b^{\text{crit}} = \frac{dPr^{G,*} - (1-d)k}{1-d+dP}.$$

Proof. Discount the first H age-specific public flows from the control date, then sum the geometric mature tail, proving (194). Substitute this expression in (193). When $H = 0$ the initial sum is empty and $d^H = 1$, including the case $d = 0$; the formulas therefore retain the current control-date flow. The mature-site expression follows directly. □

The public flows include actual household support and public bills, including any early deficits. Parent-paid installation and adoption costs enter private participation and are not deducted a second time from public income. This formula treats a stationary funded age profile, not an arbitrary stochastic policy with changing laws after control. On a changing normalized path, the exact reduced-model test instead is

$$b_t < \max_{\tau \in \mathcal{T}_{t+1}} \{-k_t(\tau) + d_t P_t(\tau)[u_{t+1}(\tau) - v_{t+1}]\}.$$

This retains the future uncontrolled value. The model provides no general primitive-only cutoff in A_C, and no stationary formula eliminates the balance state when finance binds.

Part V

Preserved 29 October comparison model

L Market fragmentation and industry change

Section 3 explains how recorded origin can become a credible condition of market access. Certification segments mineral supply and allows the North to direct its demand toward certified production while the South continues to buy from both markets. These trading opportunities determine the receipts that private producers can earn and the mining capacity they have reason to develop. I study the macroeconomic effects of this segmentation, beginning with the immediate responses of world prices, production and producer incomes.

I then study a Schumpeterian growth mechanism. Expected gains from an owner's production can make costly, risky research worth financing. Successful prototypes improve machines or expand productive installations, and disclosure carries those advances into the standard available for later investment and research. Directed technical change asks whether improving clean or dirty production offers the more attractive investment. North–South demand and certified access shape that comparison through the prices, costs and markets an installation faces. The inherited technology state shapes the initial research return, and each research round changes the markets facing the next. I derive when clean research can take over privately, the payment needed to support that allocation, and the conditions under which support can later end. These repeated improvements change industry growth and the legal tax opportunity behind governments' decisions in Section M.

L.1 World trading opportunities

Gold already reaches buyers in several large markets. Table 14 shows the changing recorded shares of mainland China and the Gulf Cooperation Council. These destinations give context to the alternative trading opportunities represented in the model.

Table 14: Recorded destinations of unwrought-gold imports.

Destination	Window	Share at start (%)	Share at end (%)
Mainland China	2015–2024	25.78	22.52
Gulf Cooperation Council	2015–2023	8.88	18.09

Notes: HS 710812. Each share is the destination's recorded import value from the 'World' partner category, divided by the corresponding total across all available reporting areas returned by the archived query for that product and year. The GCC contains Bahrain, Kuwait, Oman, Qatar, Saudi Arabia and the United Arab Emirates. Its 2024 total is unavailable because the UAE observation is missing. Recorded destinations show alternative markets; they are not counts of independent buyers or effects of introducing certification. *Source:* Frozen UN Comtrade destination-import archive used in the 4 October 2026 empirical companion, with original API queries retained.

Credible eligibility changes which of these trading opportunities a producer can use. Prices and production then adjust together. To determine that response, the market must retain the regional production choices behind supply and the actual contracts under which buyers purchase it.

L.2 North and South as buying markets

The North, $\mathcal{N}$, is the buying bloc subject to the certified-access restriction. The South, $\mathcal{S}$, can buy certified and outside mineral. These labels describe the reach of buyer regulation, rather than latitude or an income classification. They are distinct from clean and dirty production, $j \in \{C, D\}$, and from the locations of mines, indexed by r. A clean mine in the South can serve buyers in both blocs. Its location does not exclude it from the certified market.

The distinction survives when certification is imperfect: North buys the certified contract, whose physical composition depends on admission and discovery. For the log-demand contract markets below, write the certified expenditure coefficient as the sum of North and South demand. Assign outside demand and the current processed-product demand to the South:

$$E_C = I_N^C + I_S^C, \qquad E_D = I_S^D, \qquad E_R = I_S^R \tag{196}$$

North has a log preference term for its certified consumption; South has separate log terms for certified, outside and processed consumption. Their positive coefficients are the I_b^j above. Both blocs face the same generalized certified cost. The additional shared-demand term introduced in Section L.7 belongs to South. This is a geographic implementation of the existing aggregate demand equations, placing their unclassified retail final demand in South. It is not a geographic description of all observed processor trade. Buyer resources fund these expenditures and leave positive consumption of the outside numeraire.

The faithful North–South specialization below instead lets South allocate a single mineral budget between the two sources. I first study pooling and substitution, then introduce partial fraud and whole-buyer selling.

L.3 Regional supply and the income it creates

A productive installation combines equipment with the inputs needed to run it. In the implemented specialization behind the supply formulas below, inherited machinery of productivity $a(A_r) > 0$ uses a purchased input u_r to produce mineral output q_r:

$$q_r = a(A_r)u_r^{\alpha_m}, \qquad 0 < \alpha_m < 1 \tag{197}$$

Here A_r indexes installed productive capacity, u_r is a physical operating input, and α_m is its output elasticity. In this specialization the gate cannot subsidize the input: it posts $p_{ur} \geq c_{ur} > 0$, sells necessities at resource cost and has no further regional cost on this branch. Household consumption opportunities are fixed, and owned outside resources finance a strictly surviving consumption plan, including at zero mining output. The household owns the mining operation and chooses its input at the gate's positive mineral quote θ_{Br}. Its net operating income is $\theta_{Br}q_r - p_{ur}u_r$, so strict concavity gives $\theta_{Br}q_r = p_{ur}u_r/\alpha_m$. The gate's net procurement bill subtracts its input margin from the mineral payment. At a fixed positive target output that bill increases with p_{ur}, so its allowed minimum is attained at $p_{ur} = c_{ur}$. The resulting implementation cost includes both the real input bill and household operating rent:

$$\mathcal{C}_r(q_r; A_r) = \frac{c_{ur}}{\alpha_m}a(A_r)^{-1/\alpha_m}q_r^{1/\alpha_m} = \frac{\kappa_r}{1+1/\eta}A_r^{-1/\eta}q_r^{1+1/\eta} \tag{198}$$

The second expression uses $\alpha_m = \eta/(1+\eta)$ and the existing normalization $a(A_r) = [c_{ur}/(\kappa_r\alpha_m^2)]^{\alpha_m}A_r^{1-\alpha_m}$. Better installed equipment lowers the inputs and procurement cost required for a given output. This specialization has no variable household mining labor. The general Selva economy retains its household labor and participation choices.

A region sells the resulting mineral under its actual contract. At a common net receipt r_g per unit, it chooses q_r to maximize $r_g q_r - \mathcal{C}_r(q_r; A_r)$. Equating the marginal implementation cost to r_g, or choosing zero when the receipt is nonpositive, gives

$$q_r = \beta_r[r_g]_+^{\eta}, \qquad \Pi_r = \frac{\beta_r[r_g]_+^{1+\eta}}{1+\eta}, \qquad \beta_r = A_r\kappa_r^{-\eta} \tag{199}$$

Thus $\eta > 0$ is the supply elasticity, β_r measures effective capacity, and Π_r is the regional seller's optimized operating profit. Regions with the same origin and contract form group g. If its finite positive effective capacity is $B_g = \int \beta_r \, d\mu_g^{\text{reg}}$, integration gives

$$Q_g = B_g[r_g]_+^{\eta}, \qquad \Pi_g = \frac{B_g[r_g]_+^{1+\eta}}{1+\eta}, \qquad q_r(Q_g) = \frac{\beta_r}{B_g}Q_g \tag{200}$$

Output is allocated in proportion to effective capacity. Under the stated measurable-selection and resource conditions, the regional allocation returns to feasible household

choices. Geography enters through the capacities and mass of regions (see Online Appendix N.1).

Different contract groups can receive different shares ω_g of an origin's gross price p_j. Their supply and profit then use different moments:

$$Q_j = p_j^{\eta} \sum_{g \in j} B_g \omega_g^{\eta}, \qquad \Pi_j = \frac{p_j^{1+\eta}}{1+\eta} \sum_{g \in j} B_g \omega_g^{1+\eta} \tag{201}$$

A higher receipt share raises both output and the return on each unit. Supply alone therefore does not determine the income available to an inventor or government; their ownership or fiscal claims also matter. Private implementation cost includes productive inputs and household operating income. Ownership, outside wealth and subsistence needs determine how that income becomes consumption. Exit requires a financed surviving household allocation, and environmental outcomes require the corresponding physical account.

These formulas apply to the common linear-receipt branch. A local nonlinear selling return must remain inside the region's production problem. Its optimized choice is solved locally and then integrated; replacing those choices with a single capacity index would lose the distribution of local market power.

L.4 The North–South economy: pooling and certification

I first consider faithful certification with no false admission, loss or storage. All eligible legal output is serviced and required to certify when sold; dirty output remains outside. Regional continuation is linear at the applicable producer quote, so the supply class just derived applies. The outside-good endowment pays for mineral spending and the stated production, verification and investment costs, leaving a strictly positive numeraire remainder.

Buyer bloc $b \in \{N, S\}$ has preferences $n_b + I_b \log M_b$, with $I_N, I_S > 0$. North values eligible service, $M_N(c, d) = c$. South values both sources through

$$M_S(c,d) = \left[\alpha^{1/\varepsilon_{NS}} c^{(\varepsilon_{NS}-1)/\varepsilon_{NS}} + (1-\alpha)^{1/\varepsilon_{NS}} d^{(\varepsilon_{NS}-1)/\varepsilon_{NS}}\right]^{\varepsilon_{NS}/(\varepsilon_{NS}-1)} \tag{202}$$
$$0 < \alpha < 1, \qquad \varepsilon_{NS} > 1$$

Here n_b is outside-good consumption, and c, d are physical clean and dirty consumption. The weight α governs South's relative preference for clean service, and ε_{NS} is its elasticity of substitution between the two sources. Homogeneity and the positive numeraire remainder imply mineral spending I_b. These are buyer expenditure weights, distinct from production labor and machine quantities. Let $I = I_N + I_S$ and $\nu = I_N / I$.

Before verifiable origin, every purchase uses an anonymous pooling facility. The pool pays the same quote p_0 for either source and delivers the common clean proportion π_{NS}. Buyers cannot order a different mix or verify private source labels. Purchasing m units therefore

yields $(c, d) = (\pi_{NS}m, (1-\pi_{NS})m)$. North receives eligible service $\pi_{NS}m$, while South receives $M_S(\pi_{NS}, 1-\pi_{NS})m$. Both spend I_b on the pool: the composition changes utility but not log-demand expenditure. With positive pre-fee capacities B_C, B_D, pooling clears at

$$p_0 = \left(\frac{I}{B_C + B_D}\right)^{1/(1+\eta)}, \qquad Q_j^0 = B_j p_0^\eta, \qquad \pi_{NS} = \frac{B_C}{B_C + B_D} \tag{203}$$

Pooling is thus a transaction restriction. Certification makes source-specific buying feasible and allows North to direct its spending to eligible output.

For this specialization, the certified buyer pays gross price p_C, and an ad valorem fee $\varphi_{NS} \in [0, 1)$ leaves the producer $(1-\varphi_{NS})p_C$. This fee convention is specific to the present CES comparison. Define $a_C = (1-\varphi_{NS})^\eta B_C$ and $a_D = B_D$, so supply is $Q_C = a_C p_C^\eta$, $Q_D = a_D p_D^\eta$. North demands I_N/p_C; South's demands are

$$(c_S, d_S) = \frac{I_S}{\alpha p_C^{1-\varepsilon_{NS}} + (1-\alpha)p_D^{1-\varepsilon_{NS}}} \left(\alpha p_C^{-\varepsilon_{NS}}, (1-\alpha)p_D^{-\varepsilon_{NS}}\right) \tag{204}$$

Market clearing requires $a_C p_C^\eta = I_N/p_C + c_S$ and $a_D p_D^\eta = d_S$. Put $z = p_C/p_D$ and $\omega_{NS} = \alpha/(1-\alpha)$. The two markets clear at the unique positive root and price level

$$\begin{aligned} \Upsilon_{NS}(z) &= (1-\nu)\frac{a_C}{a_D} z^{1+\eta} - \omega_{NS} z^{1-\varepsilon_{NS}} - \nu = 0 \\ p_D^{1+\eta} &= \frac{I}{a_D + a_C z^{1+\eta}}, \qquad p_C = z p_D \end{aligned} \tag{205}$$

The root is unique because Υ_{NS} is strictly increasing from minus infinity to plus infinity. Clean demand combines North's exclusive certified spending with South's clean purchases; outside dirty demand comes only from South.

L.4.1 The price shock after certification

Before certification, both sources sell at p_0, financed by pooled North and South demand. After certification, North directs its spending to clean supply, while South chooses between the two sources at p_C, p_D. The net price effect depends on North's expenditure share, South's preferences, productive capacities and certification costs.

Proposition L.1 (The certification price shock)**.** *Certification lowers the outside price below the pooled price exactly when Northern demand and Southern preference for certified mineral are large relative to clean capacity. Hold technologies and buyer budgets fixed, before any innovation occurs. Then $p_D < p_0$ if and only if*

$$(1-\nu)\frac{B_C}{B_D} < \nu + \omega_{NS}(1-\varphi_{NS})^{\eta(\varepsilon_{NS}-1)/(1+\eta)} \tag{206}$$

Under this condition the gross clean buyer quote rises,

$$\frac{p_C}{p_0} > (1-\varphi_{NS})^{-\eta/(1+\eta)} \geq 1 \tag{207}$$

and for dirty regions with unchanged productive capacity and selling contracts, output, gross receipts and optimized operating profit fall. With superscripts 0 and 1 denoting pooling and certification,

$$\begin{aligned} \frac{q_r^1}{q_r^0} &= \left(\frac{p_D}{p_0}\right)^{\eta} \\ \frac{p_D q_r^1}{p_0 q_r^0} &= \frac{\Pi_r^1}{\Pi_r^0} = \left(\frac{p_D}{p_0}\right)^{1+\eta} \end{aligned} \tag{208}$$

At equality in (206) the dirty quote is unchanged; reversing the inequality raises it (see Appendix Q.7.1).

Scope. The clean producer receives $(1-\varphi_{NS})p_C$, and its net quote rises only if $(1-\varphi_{NS})p_C > p_0$. South buys positive dirty output at every finite positive capacity state.

Receipts and profit have the same proportional change because optimized profit is the fixed fraction $1/(1+\eta)$ of receipts. Sections 2–3 studied buyer pressure while remaining selling opportunities could absorb production. This price shock changes the outside quote itself, reaching the return on every dirty unit in this linear-continuation class.

L.4.2 Segmentation and the global policy ceiling

Theorem 2.16 holds unaffected demand and the value of laundering fixed. Proposition L.1 moves the outside market in which both are earned. To the extent that unaffected bids and laundering proceeds are priced at the outside quote, condition (206) is the condition under which segmentation loosens that ceiling. A formal statement requires embedding the VCG cell of Section L.6 in the outside market.

L.4.3 Market size and the return to clean capacity

The resulting gross source expenditures and producer operating surpluses are

$$\begin{aligned} E_D^{NS} &= \frac{I_S}{1+\omega_{NS} z^{1-\varepsilon_{NS}}}, \quad E_C^{NS} = I - E_D^{NS} \\ I_N \le E_C^{NS} &\le I, \qquad\qquad 0 \le E_D^{NS} \le I_S \\ \Pi_C^{NS} &= \frac{(1-\varphi_{NS})E_C^{NS}}{1+\eta}, \quad \Pi_D^{NS} = \frac{E_D^{NS}}{1+\eta} \end{aligned} \tag{209}$$

North supplies a certified revenue floor, while dirty producers compete for at most South's mineral expenditure. At fixed budgets and fee, greater relative clean capacity a_C/a_D lowers z. Because $\varepsilon_{NS} > 1$, South substitutes toward clean mineral: E_C^{NS} rises and E_D^{NS} falls. With dirty capacity fixed, both unit prices fall even as total clean operating surplus rises. These are the price and market-size channels through which productive improvements change the return available to owners.

Lemma L.2 (Segmentation and the research prizes)**.** *Relative to pooling, certification lowers dirty producers' income exactly under (206), and raises clean producers' net income*

exactly when their net spending share exceeds their pooled share. Let $E_j^0 = IB_j/(B_C + B_D)$ *be pooled spending on source* j*, and* E_C^{NS}, E_D^{NS} *the certified-market spending in* (209)*. Then*

1. $E_D^{NS} < E_D^0$ *if and only if* (206) *holds, so the dirty operating surplus* $\Pi_D^{NS} = E_D^{NS}/(1+\eta)$ *falls below its pooled value exactly under* (206)*;*

2. *the clean operating surplus* $\Pi_C^{NS} = (1-\varphi_{NS})E_C^{NS}/(1+\eta)$ *exceeds its pooled value* $E_C^0/(1+\eta)$ *if and only if* $(1-\varphi_{NS})E_C^{NS} > E_C^0$*. With* $\varphi_{NS} = 0$ *the two statements are equivalent.*

Proof. By (203), pooled receipts of source j are $p_0 Q_j^0 = B_j p_0^{1+\eta} = IB_j/(B_C + B_D) = E_j^0$. Under certification $a_D = B_D$, so $E_D^{NS} = a_D p_D^{1+\eta} = B_D p_D^{1+\eta}$, and $E_D^{NS} < E_D^0$ if and only if $p_D < p_0$, which is (206) by Proposition L.1. On the linear-receipt branch optimized profit is the share $1/(1+\eta)$ of receipts (199); the clean producer receives the share $1-\varphi_{NS}$ of gross spending (209). Both regimes spend the whole budget I on mineral, so $E_C^{NS} + E_D^{NS} = E_C^0 + E_D^0 = I$. This gives part 2, and with $\varphi_{NS} = 0$ it makes $E_C^{NS} > E_C^0$ equivalent to $E_D^{NS} < E_D^0$. □

This faithful CES economy isolates the North–South substitution mechanism. The contract markets that follow keep the geographic access distinction while allowing imperfect certification and nonlinear selling returns. They use the log-demand coefficients in (196); their prices and owner incomes must be solved under those contracts. The CES comparative statics require the preferences, fee incidence and faithful linear-continuation conditions used here.

L.5 Certified and outside prices

Now allow false admission and buyer diligence under the public-quote institution. A buyer purchases one physical unit and its origin record, accepting a public lottery over the origins actually supplying certified trade. After assignment, the buyer observes its origin and contract and chooses diligence. Clean suppliers receive $h > 0$; dirty suppliers receive $p_D = xh$, with $0 < x < 1$. A separate licensed intermediary organizes admissions and owns the return from any false admission. Its profit is not automatically the producer's or research owner's income.

At each vulnerable origin k, the Section 4 contract supplies a late-discovery probability r_k, distinct collectible bonds hb_k, early and late buyer losses hl_{Ek} and hl_{Lk}, diligence cost $h\kappa_k e^2/2$, and forgery cost h times a uniform draw on $[0, c_k]$. The reporting incentive, accessible evidence and separate bribe liquidity remain in force. On the interior branch,

$$a_k = \frac{r_k l_{Lk} - l_{Ek}}{\kappa_k}, \qquad \pi_k(x) = \frac{1 - x - r_k b_k}{c_k + b_k(1-r_k)a_k}, \qquad e_k = a_k \pi_k \tag{210}$$

Here π_k is the equilibrium share of false admissions and e_k is buyer diligence. The contract requires $r_k l_{Lk} - l_{Ek} > (1-r_k)b_k$, positive c_k, κ_k, and nonnegative early loss. An origin

without illegal stock has no illegal-origin admission action. A larger certified–dirty price difference raises the gain from fraud; investigation and collectible sanctions lower it. Buyers investigate more when false admission is more likely. Those choices meet at the equilibrium price ratio.

Let B_{Ck} and B_{Dk} be clean and dirty regional supply coefficients. Production and its two destinations obey

$$\begin{aligned} Q_{Ck} &= B_{Ck}h^{\eta}, \quad Q_{Dk} = B_{Dk}(xh)^{\eta} \\ X_k &= \frac{B_{Ck}h^{\eta}}{1-\pi_k}, \qquad Y_k = Q_{Dk} - \pi_k X_k \end{aligned} \tag{211}$$

Every false admission uses actual illegal-origin material. The selected branch requires $Y_k > 0$ at every vulnerable origin as well as interior fraud and diligence. With $X = \sum_k X_k$ and $Y = \sum_k Y_k$, total extraction is $X + Y$. Subsequent withdrawal changes the live status of part of X; it does not undo its extraction. Dirty suppliers receive p_D for material entering either destination.

This supply calculation embeds delivered mineral in durable products and excludes recycling or return to current primary supply. Its origin evidence remains in the state. In the wider institutional application, processing and recycling inherit recorded history without fresh clean reminting. That continuing history is compatible with the primary-supply subclass; adding recycled deliveries to its current flow would require a release and return account.

Buyers pay for the certified unit and its credibility. This market uses contracts scaled by the clean quote h. Section 4's fee ϕ was a currency amount; here ϕ denotes a dimensionless rate, so the service fee per unit is $h\phi$. Let $R_k(x)$ be the buyer's diligence and retained-loss cost at origin k, divided by h. The generalized certified unit cost is $hK(x)$, where

$$K(x) = 1 + \phi + \sum_k \frac{X_k}{X} R_k(x)$$

Only the invoice $h(1+\phi)X$ reaches suppliers and service providers; the buyer also bears its diligence and retained losses. Quasilinear log demands give $hKX = E_C$ and $xhY = E_D$, with positive expenditure coefficients and sufficient buyer liquidity. North and South together finance certified demand, while South finances outside demand as specified in (196).

Put $B_D = \sum_k B_{Dk}$, and define

$$\begin{aligned} F_{\mathrm{cw}}(x) &= \sum_k \frac{B_{Ck}}{1-\pi_k(x)}, & H_{\mathrm{cw}}(x) &= F_{\mathrm{cw}}(x) - \sum_k B_{Ck} \\ T(x) &= \sum_k \frac{B_{Ck}[1+\phi+R_k(x)]}{1-\pi_k(x)}, & J(x) &= \frac{x[B_D x^{\eta} - H_{\mathrm{cw}}(x)]}{T(x)} \end{aligned} \tag{212}$$

Theorem L.3 (Common-price equilibrium with fraud and diligence)**.** *Certified and outside prices, production, admissions and contamination are determined jointly and uniquely on each contract branch. Take a closed interval strictly inside* $(0,1)$ *on which the individual incentive, funding, liquidity and origin-stock conditions hold, without fraud or effort*

switching branches. If its endpoint values of J straddle E_D/E_C, the equilibrium in this contract branch is unique and satisfies

$$J(x^*) = E_D/E_C, \qquad h^* = \left[\frac{E_C}{T(x^*)}\right]^{1/(1+\eta)}, \qquad p_D^* = x^* h^* \tag{213}$$

On compact strictly feasible primitives the selected prices and values are continuously differentiable with bounded derivatives.

The result determines production, admissions, composition and surviving claims together. At a fixed clean quote, a higher dirty price expands dirty production and reduces diversion into certification. More material remains for outside trade, while lower fraud reduces the certified contract's expected diligence and loss cost. These forces make J increasing; clean clearing then fixes the price level (see Online Appendix N.2).

Certification thus creates distinct prices and investment opportunities while illegal material can still be sold. Lower surviving contamination and lower illegal extraction are separate outcomes, connected through these prices and the regional production response. The discovery comparison in Section 4 concerns the former within certification.

L.6 Whole-buyer selling and processing access

Public quotes describe one selling institution. The finite whole-buyer Vickrey–Clarke–Groves (VCG) institution instead makes the producer's return nonlinear in inventory. Each region contains an atomless mass of independently operated cells, each serving its own finite buyers. A processor of productivity v converts one mineral unit into v final-product units sold at the common price p.

Processors have exclusive, enforceable access to their cell's mineral: they cannot obtain an equivalent public input, create extra identities or enter another cell. Producers have access only to assigned processors and the genuine retained-conversion route. Refusing the contract permits exit at zero; it does not open the bulk raw-mineral market. These access restrictions are fixed before participation and procurement. They are maintained market assumptions, supported by the specified physical custody arrangement rather than implied by a ledger alone.

The consignment contract divides regional stock equally across installed cells. If n_r is their mass, the region's continuation is $W_r(Q;p) = n_r W^V(Q/n_r;p)$. Unsold mineral uses a separate genuine conversion technology; it cannot later be released to primary buyers. Each extracted unit enters primary delivery or retained conversion, while terminal trade reallocates already delivered material (see Online Appendix O.3).

For finitely many cell types, actual final supply is $Z_R(p) = \sum_j n_j Z_j(p)$, and log final demand clears at

$$p \sum_j n_j Z_j(p) = E_R \tag{214}$$

Final output follows the actual allocation and conversion; it is not receipts divided by p. The strict cell branches have $Z_j'(p) \geq 0$. If the endpoint signs enclose a root and a positive core remains, the selected root in that interval is unique and continuous in cell masses and expenditure. Regional production and profit are $n_r q_j(p)$ and $n_r \pi_j(p)$, with attained household implementations. An innovation that replicates a cell adds its equipment, authorized processing access and institutions together.

L.7 Price feedback and the opportunity to invest

Bulk and retail routes use disjoint origin and contract groups, so each physical mineral unit enters one route. They can still share final consumers. In the geographic demand implementation, South supplies this shared consumption: the linked preference is added to its separate log terms, while North retains certified demand. With the additional preference term $\varepsilon \log(Y + \gamma_Z Z_R)$, $\varepsilon, \gamma_Z > 0$, their prices satisfy

$$\begin{aligned} hK(x)X &= E_C \\ hxY &= E_D + \frac{\varepsilon Y}{Y + \gamma_Z Z_R(p)} \\ pZ_R(p) &= E_R + \frac{\varepsilon \gamma_Z Z_R(p)}{Y + \gamma_Z Z_R(p)} \end{aligned} \tag{215}$$

Processed supply now changes demand for outside mineral, and outside supply changes demand for the processed product. Reform or research in one set of origins can therefore change the prices and actual producer incomes facing the other. The connected stage requires a common compact domain with positive deliveries, funded contracts, strict cell choices and a nonsingular price system. A growing economy additionally requires that this domain remain valid along its path. A single date's VCG solution does not establish those infinite-path conditions.

The connected growth constructions retain their own stage assumptions. In the mixed-scale construction, safe supply and retail installations expand while vulnerable capacity, outside supply and outside expenditure retain positive absolute cores. A falling dirty share therefore need not eliminate absolute illegal extraction. The financed faithful retail subclass also preserves a separate vulnerable bulk group: retail custody uses knowing bonded participants with veto power, an unsuppressible late audit and distinct collectible pools. A nonmineral appropriation pays installation, capital carry, service, permanent evidence and reserves.

Suppliers receive their actual contract receipts; organizers keep their separate fraud surplus; service providers receive their funded fees. Aggregate innovation and government reform change the common prices at which these incomes are earned. Research responds to the resulting income available to an owner from a successful project.

L.8 Machines, production and the return to an improvement

Consider the owner of a clean mine. It pays for equipment, operating inputs and certified access, then earns receipts from the mineral it sells to eligible buyers in North and South. An improved machine can make those inputs more productive; a new complete installation can add capacity. A dirty operation offers a second possible use of research. Investors compare the returns they can earn from these improvements, and scientific effort changes the productive technologies available in each industry.

A machine-production benchmark shows how equipment generates this return. As in the directed technical change framework of Acemoglu [2002] and Acemoglu et al. [2012], sector $j \in \{C, D\}$ combines labor ℓ_j with a unit mass of machine lines $m \in [0, 1]$:

$$Y_j = \frac{\ell_j^{\beta_m}}{1-\beta_m} \int_0^1 \mathcal{A}_j(m)^{\beta_m} x_j(m)^{1-\beta_m}\, dm, \qquad 0 < \beta_m < 1 \tag{216}$$

Here Y_j is physical sector output, $x_j(m)$ is the quantity of machine service on line m, and $\mathcal{A}_j(m)$ is that line's quality. The parameter β_m is the labor exponent, distinct from regional effective capacity β_r. At sector output price θ_j, a competitive producer equates the price of each machine to its marginal contribution to receipts. Its inverse demand is

$$p_{x,j}(m) = \theta_j \mathcal{A}_j(m)^{\beta_m} x_j(m)^{-\beta_m} \ell_j^{\beta_m} \tag{217}$$

In this machine-production benchmark, an exclusive supplier pays $1-\beta_m$ units of the numeraire per machine unit. Maximizing $[p_{x,j}(m) - (1-\beta_m)]x_j(m)$ gives the monopoly price $p_{x,j}(m) = 1$. Let $\overline{\mathcal{A}}_j = \int_0^1 \mathcal{A}_j(m)\, dm$ be average machine quality. Substitution gives machine demand, its supplier's operating profit, and sector output:

$$\begin{aligned} x_j(m) &= \mathcal{A}_j(m)\theta_j^{1/\beta_m}\ell_j \\ \pi_j^M(m) &= \beta_m \mathcal{A}_j(m)\theta_j^{1/\beta_m}\ell_j \\ Y_j &= \frac{\overline{\mathcal{A}}_j \ell_j}{1-\beta_m}\theta_j^{(1-\beta_m)/\beta_m} \end{aligned} \tag{218}$$

Machine demand equals line quality times employment and the output-price term; at the normalized machine price of one, supplier profit is the fraction β_m of that demand. The last expression aggregates production across machine lines using average quality. An improvement's value therefore depends on its quality, the output price and the scale at which it can be used. The machine supplier earns π_j^M under exclusive supply rights; this is the patent-based benchmark. In the regional economy studied here, the investor instead receives the return on its own productive installation or an explicitly contracted claim. Disclosed designs can be copied without a royalty. An armed group's control of a mine can give it operating income without giving it a patent on the machinery's design.

The North–South market-size channel is the expenditure that an adopting installation can serve. Its return is realized through the equilibrium output price and its actual

productive scale, including the employment ℓ_j used in this benchmark. The installed-equipment production and cost formulas above give the corresponding owned operating return directly. At the actual net receipt r_i, an owner with baseline effective capacity β_i^0 earns $\Pi_i^0 = \beta_i^0[r_i]_+^{1+\eta}/(1+\eta)$. If a successful prototype raises its own capacity by the proportion $\gamma_i > 0$, reoptimizing production at the same public prices gives

$$\Delta q_i = \gamma_i \beta_i^0 [r_i]_+^{\eta}, \qquad \Delta\Pi_i = \frac{\gamma_i \beta_i^0 [r_i]_+^{1+\eta}}{1+\eta} = \gamma_i \Pi_i^0 \tag{219}$$

For example, a clean owner selling at the certified quote h uses its actual net receipt after the contract's deductions in r_i. A dirty owner's receipt comes from its own selling opportunity. At positive receipts, a larger price or a larger owned capacity raises the operating gain in this branch. When receipt shares differ, the profit moment derived above determines the gain belonging to each owner. The origin's total sales cannot replace that ownership calculation.

The same logic values replication in the VCG economy. Adding an independent complete installation adds its optimized cell profit, with its equipment, finite buyers, authorized processing access and retained-conversion route. The value includes the auction and retention decisions already solved above. Improving costs inside one nonlinear installation requires solving its new production problem. Thus the machine calculation explains the price and scale channels, while the actual regional and VCG operating values determine the research return used below. World prices continue to satisfy the preceding market equations.

L.9 From production to the investor's innovation value

The Schumpeterian cycle proceeds through repeated research, deployment and diffusion. At date t, a project owner chooses costly scientific effort in anticipation of the net gain it can keep from a successful improvement. A successful prototype first operates at $t+1$, giving the owner a one-interval productive lead. Disclosure then makes the improved standard available for the next installation and research round. Designs can be copied without a royalty; equipment, site adaptation, custody and market access still require paid implementation. Production income arrives after research and installation have been paid for. The investor's gain comes from its own mine, productive portfolio or an explicitly contracted claim, valued with its continuing costs and ownership.

Let $\mathcal{V}_{i,t}^{\mathrm{s}}(\Omega)$ and $\mathcal{V}_{i,t}^{\mathrm{f}}(\Omega)$ denote the date-t expected discounted value of the owner's future net cash flows after success and failure. Both use the same actual equilibrium prices and anticipated public market path Ω: a single owner has zero mass and takes that path as given. They include operating income, installation and access costs, collateral, later payments and refunds, and residual ownership value. Current scientific wages, laboratory costs and the separate research success payment are excluded. The incremental private innovation value is

$$J_{i,t} = \mathcal{V}^{\mathrm{s}}_{i,t}(\Omega) - \mathcal{V}^{\mathrm{f}}_{i,t}(\Omega), \qquad \widehat{J}_{i,t} = J_{i,t}/c_t \tag{220}$$

The scale is $c_t = 1$ in the fixed-expenditure benchmark and $c_t = L_t = g^t$ in the growing-income constructions. Consequently $J_{i,t}$ is a date-t value in numeraire units, and $\widehat{J}_{i,t}$ is its value per unit of the common income scale.

To compute it, let Δ mean success minus failure. Write $\mathcal{O}$ for owned operating receipts net of current production and service costs, K for the installation and access bill, and B for a new collateral deposit. Let $C_{i,t+u}$ record subsequent net payments from the affected vintage, with refunds entered as negative payments. For the fixed finite contractual claims horizon T_H, the complete value is

$$\begin{aligned} J_{i,t} = \delta_R \mathbb{E}[\Delta \mathcal{O}_{i,t+1} - \Delta K_i - \Delta B_i] - \sum_{u=2}^{T_H+1} \delta_R^u \mathbb{E}\Delta C_{i,t+u} \\ + \delta_R^{T_H+1} \mathbb{E}\Delta V^{\mathrm{other}}_{i,t+T_H+1} \end{aligned} \tag{221}$$

Here $\delta_R \in (0,1)$ is the owner's discount factor. The first term discounts the date-$t+1$ operating gain and installation and collateral payments back to date t. The sum retains the later payments and refunds. Production begins at $t+1$, so the prototype's liabilities close by $t + T_H + 1$. The final term retains any remaining difference in wealth, durable machines, ownership or operating opportunities. Original stock, committed selling terms and evidence stay in the affected vintage's state.

This expression links the productive improvement to the research incentive. If the investor owns the regional seller's payoff on the linear capacity branch, and all current service deductions are included in r_i or cancel between success and failure, then $\Delta\mathcal{O} = \gamma_i \Pi_i^0$. The household operating rent is already paid through the implementation cost. Any further current payment borne by the investor remains in the actual $\Delta\mathcal{O}$. For independent complete VCG installations, use the corresponding increment in owned operating value, including their actual current costs. The other cash-flow terms then determine how much of this operating improvement the research owner can keep.

A one-interval production lead need not make the innovation value one-period. Its residual ownership term cancels only under the stated one-interval lease or liquidation arrangements and identical later opportunities. Later vintage payments remain until settled. If the owner posts its own collateral, the discounted deposit and refund already value its temporary loss of access to that capital. A facility-financed deposit instead uses the actual loan payments. These arrangements preserve the distinction between owning a productive asset, paying for an installation, and owning an enforceable claim on someone else's receipts.

L.10 Schumpeterian innovation and the direction of research

Research makes the productive improvement more likely. A type-i project employing s scientists succeeds with probability $h_i(s) = 1 - e^{-\nu_i^R s}$, where $\nu_i^R > 0$ measures research

productivity. The convex nondecreasing laboratory cost is $c_t R_i(s)$, with $R_i(0) = 0$. One scientist costs $c_t w_t$ at date t, where $w_t \geq 0$ is the common normalized wage. A funded success payment has date-t value $c_t G_i \geq 0$ and is actually paid at $t+1$ as $c_t G_i/\delta_R$.

The owner compares the expected innovation value with these research costs. After subtracting the failure value, which does not vary with effort, and dividing by c_t, its choice is

$$s_{i,t} \in \arg\max_{s \geq 0} \left\{ h_i(s)[\widehat{J}_{i,t} + G_i] - R_i(s) - w_t s \right\} \tag{222}$$

Zero effort preserves the failure continuation value. At positive effort the owner equates the marginal expected innovation value, net of laboratory cost, to the scientific wage. Prices and market access matter because they enter the owned production gain inside $J_{i,t}$; the success probability translates that gain into the amount the investor is willing to pay for research.

Scientific labor must also be feasible. I reserve $S_V \in (0, S)$ scientists for paid verification research. If N_i is the mass of research owners of type i, distinct from regions and buyers, then

$$S_V + \sum_i N_i s_{i,t} \leq S \tag{223}$$

I choose a clean research core $*$, with $s_* = (S - S_V)/N_* > 0$. Let $\underline{J}_i$ and $\overline{J}_i$ bound the complete normalized innovation value $\widehat{J}_{i,t}$ over the stated market domain. Fully private mineral research sets every relevant G_i to zero. It directs all available mineral scientists to the clean core if that core's marginal return at s_* exceeds both every competing project's initial marginal return and zero. With nonnegative success payments the same sufficient condition is

$$h'_*(s_*)[\underline{J}_* + G_*] - R'_*(s_*) > \max\left\{0, \max_{i \neq *} \left[h'_i(0)[\overline{J}_i + G_i]_+ - R'_i(0)\right]\right\} \tag{224}$$

The bounds cover actual prices and all continuing private obligations. The strict comparison makes moving scientific effort from a competing project, or from unused labor, to the clean core profitable. The competitive allocation therefore has $N_* s_* = S - S_V$. Independent successes and the specified aggregate research technology give the clean standard the growth factor

$$g = 1 + \gamma_* h_*(s_*) > 1 \tag{225}$$

The technology growth rate $g - 1$ is the innovation success probability $h_*(s_*)$ times the improvement step γ_*. Under the stated aggregate research technology, repeated research rounds raise the clean standard by the factor g. The selected allocation produces an expanding clean technology while the other mineral research types receive no scientific

effort. This is a sufficient condition throughout the stated market domain. Certification changes the returns being compared; whether they induce this allocation depends on the owned production opportunities, research costs and funding. Failure of the conservative bound does not establish that private research fails at actual equilibrium prices, or that partial clean research cannot support a transition.

L.11 Path dependence, subsidy thresholds and private takeover

A research owner inherits productive capacity and a market in which to use it. Its investment changes the standard available in the next round, while aggregate innovation changes the prices and spending from which later owners earn income. This feedback creates path dependence: an initial clean lead or an early period of clean research can change subsequent research incentives. Along the clean-directed CES path below, clean research expands the clean production market and reduces its rival's. Support can be temporary when this feedback makes private clean research attractive enough to retain all mineral scientists. By *clean takeover* I mean that all available mineral scientists choose the clean core. Partial clean research can occur before that threshold, and full clean research does not immediately eliminate dirty production.

L.11.1 The payment needed at market values

I assume differentiable laboratory costs for these marginal comparisons. I first evaluate the complete innovation values $\widehat{J}_{i,t}$ at prices and a future market path consistent with the proposed full-clean allocation s_*. I hold other current payments and anticipated future policy fixed. Ownership and financing must also be those of the proposed policy and continuation. Changing support or its continuation can change the complete values and requires recomputing the threshold.

Theorem L.4 (The research support threshold)**.** *Directing all mineral science to clean technology needs a payment exactly when the private clean prize falls short of the competing return. At prices and a market path consistent with the full-clean allocation s_*, let $\mathcal{B}_t$ be the largest competing marginal research return at zero scientific effort, including the option of leaving science unused, and define*

$$\begin{aligned} \mathcal{B}_t &= \max\left\{0, \max_{i\neq *}\left[h_i'(0)[\widehat{J}_{i,t}+G_i]_+ - R_i'(0)\right]\right\} \\ \Gamma_t &= \frac{R_*'(s_*)+\mathcal{B}_t}{h_*'(s_*)} - \widehat{J}_{*,t} \\ G_{*,t}^{\min} &= \max\{0,\Gamma_t\} \end{aligned} \tag{226}$$

A nonnegative payment weakly supports this allocation if and only if $G_{*,t} \geq G_{*,t}^{\min}$. A strict marginal advantage requires $G_{*,t} > \Gamma_t$, together with $G_{*,t} \geq 0$. Thus $\Gamma_t < 0$ means zero clean support already strictly supports full clean allocation; $\Gamma_t = 0$ gives a marginal equality

without clean support; and $\Gamma_t > 0$ *requires a positive payment for this allocation at those consistent prices.*

Scope. These are statements about the proposed allocation at consistent prices, rather than a test ruling out partial clean research or other market continuations. Changing support or its continuation changes $\widehat{J}$ and requires recomputing the threshold. Equality need not create multiple research choices.

The current clean success payment has the same normalized date-t units as G_* above. The threshold solves the clean core's marginal-return condition for the payment: the core must be able to pay a scientific wage at least as high as the competing return.

If support is a matching rate μ_t, set $G_{*,t} = \mu_t \widehat{J}_{*,t}$. For a positive private innovation value, the strict subsidy threshold is $\mu_t > \Gamma_t / \widehat{J}_{*,t}$, with $\mu_t \geq 0$. Multiplying a nonpositive private value cannot provide a positive reward; the additive payment remains available in that case. These payments retain the date-$t+1$ settlement and funding rule already specified.

A conservative policy can instead use the complete-value bounds in (224). Let $\overline{\mathcal{B}}$ be the same competing-return comparison with each $\widehat{J}_{i,t}$ replaced by its valid upper bound. Define

$$\overline{\Gamma} = \frac{R'_*(s_*) + \overline{\mathcal{B}}}{h'_*(s_*)} - \underline{J}_* \tag{227}$$

A nonnegative clean payment strictly above $\overline{\Gamma}$ is sufficient throughout that domain. Failure of this uniform sufficient test does not establish that support is needed at the actual state. Both thresholds use the full owner-specific value, including the dated liabilities and residual ownership in (221).

L.11.2 An explicit inherited-capacity threshold

The faithful North–South CES economy makes the feedback transparent. For this comparison, I take two research types, clean C and dirty D, each containing a unit mass of otherwise identical atomistic project owners, and write $\bar{S} = S - S_V$ for available mineral scientists. I keep buyer budgets, the fee, differentiable convex laboratory costs and the research productivities ν_j^R fixed. Let $a_t = a_{C,t}/a_{D,t}$ be inherited relative effective capacity. Its evolution is

$$v_j(s) = 1 + \gamma_j h_j(s), \qquad \frac{a_{t+1}}{a_t} = \frac{v_C(s_{C,t})}{v_D(s_{D,t})} \tag{228}$$

History enters through installed capacity and the market returns it creates. Under the candidate full-clean allocation, put $\bar{v}_C = v_C(\bar{S}) > 1$. The next production round then has relative capacity $\bar{v}_C a_t$; let $E = E_C^{NS}(\bar{v}_C a_t) \in (I_N, I)$ be its clean expenditure, using Section L.4's market clearing.

For a closed-form private threshold, I use a transparent one-interval subclass. Owners are risk neutral with additive outside-good wealth. Prototypes expire after one production interval. Later funded opportunities coincide after success and failure, and success has no later marginal-utility or treasury effect. All other success-minus-failure contract payments,

refunds and ownership values cancel. The constant additional installation and access bill is $K_j \geq 0$. These restrictions reduce the complete value to a temporary operating gain; they do not replace the general value calculation for other contracts or the nonlinear VCG economy. At the actual next-round prices, the two private prizes are

$$\begin{aligned} J_C(E) &= \delta_R \left[\frac{\gamma_C(1-\varphi_{NS})E}{(1+\eta)\bar{v}_C} - K_C \right] \\ J_D(E) &= \delta_R \left[\frac{\gamma_D(I-E)}{1+\eta} - K_D \right] \end{aligned} \tag{229}$$

The factor $1/\bar{v}_C$ recovers the unsuccessful clean owner's capacity from the aggregate sector after research. South's substitution toward expanding clean capacity raises E, increasing the clean prize and reducing the dirty prize. This is the path-dependent feedback from research to the next market and back to research.

To express the private cutoff, collect the fixed coefficients of the two marginal returns:

$$\begin{aligned} u_C &= \frac{h_C'(\bar{S})\delta_R\gamma_C(1-\varphi_{NS})}{(1+\eta)\bar{v}_C}, & c_C &= h_C'(\bar{S})\delta_R K_C + R_C'(\bar{S}) \\ u_D &= \frac{h_D'(0)\delta_R\gamma_D}{1+\eta}, & c_D &= h_D'(0)\delta_R K_D + R_D'(0) \end{aligned} \tag{230}$$

Here $u_C, u_D > 0$ and $c_C, c_D \geq 0$. The clean marginal return at full effort is $u_C E - c_C$; dirty research's marginal return at zero effort is $u_D(I-E) - c_D$.

Theorem L.5 (Inherited capacity and temporary support)**.** *In the one-interval CES subclass, clean research takes over privately once inherited relative clean capacity exceeds an explicit cutoff, and support can end once capacity crosses it. Without research payments, the clean return strictly exceeds both its rival and zero exactly when*

$$E > E^\dagger, \qquad E^\dagger = \max\left\{ \frac{c_C}{u_C}, \frac{u_D I + c_C - c_D}{u_C + u_D} \right\} \tag{231}$$

For $I_N < E^\dagger < I$, the CES market equations translate this expenditure cutoff into relative capacity:

$$\begin{aligned} a^\dagger &= \frac{E^\dagger}{I - E^\dagger} \left[\frac{E^\dagger - I_N}{\omega_{NS}(I - E^\dagger)} \right]^{(1+\eta)/(\varepsilon_{NS}-1)} \\ \textit{strict private takeover at date } t \quad &\Longleftrightarrow \quad a_t > a^\dagger/\bar{v}_C \end{aligned} \tag{232}$$

Three cases follow.

1. Private clean research is already viable. *If current $E > E^\dagger$, no mineral-research payment is needed to sustain the selected full-clean continuation. In particular, $E^\dagger \leq I_N$ makes the private inequality strict at every finite positive capacity state.*

2. Support can be temporary. *If $I_N < E^\dagger < I$ and the inherited state has not crossed the cutoff, a funded payment satisfying Theorem L.4 sustains full-clean research while*

capacity grows along $a_t = \bar{v}_C^t a_0$. *Payments can end at the first date* T *for which* $\bar{v}_C^{T+1} a_0 > a^\dagger$.

3. The limiting private return can be too low for finite withdrawal. *The limiting clean private prize in this fixed-budget subclass is*

$$J_C^\infty = \delta_R \left[\frac{\gamma_C (1 - \varphi_{NS}) I}{(1 + \eta) \bar{v}_C} - K_C \right] \tag{233}$$

Finite-date strict private takeover is possible exactly when $u_C I - c_C = h_C'(\bar{S}) J_C^\infty - R_C'(\bar{S}) > 0$. *If* $u_C I \leq c_C$, *no finite state supports strict private full-clean research in this subclass, and maintaining that allocation requires continuing payments.*

(see Appendix Q.9.1)

Scope. The result establishes existence and persistence of the selected private continuation, not global uniqueness, and it holds buyer budgets fixed. It does not remove financing and coordination requirements. At the cutoff the private comparison is weak rather than strict, and partial clean research can occur below it.

The division by $\bar{v}_C$ reflects timing: date-t research earns its operating return at the next date's improved aggregate capacity. Once payments end, the private inequality strengthens as clean capacity continues to grow. At $u_C I = c_C$, the required additive payment approaches zero but remains positive at finite dates; at $u_C I < c_C$, its limit is positive. Support remains feasible with adequate outside funding. The obstruction concerns private full-clean research, rather than the possibility of partial clean effort or legal production.

Corollary L.6 (Certification and private takeover)**.** *If certification raises clean producers' net spending at next-round capacities, every state with private takeover under pooling also has it under certification. Define pooled research prizes by* (229) *with* E *replaced by pooled clean spending* E_C^0 *and* $\varphi_{NS} = 0$, *at the same next-round capacities, research technologies and installation bills. If* $(1 - \varphi_{NS}) E_C^{NS} \geq E_C^0$, *strict private full-clean research under pooling implies strict private full-clean research under certification.*

Proof. Under pooling the clean marginal return at full effort is $u_C^0 E_C^0 - c_C$ with $u_C^0 = u_C/(1 - \varphi_{NS})$, and dirty research's marginal return at zero effort is $u_D(I - E_C^0) - c_D$. Under certification they are $u_C E_C^{NS} - c_C$ and $u_D(I - E_C^{NS}) - c_D$. The hypothesis gives $u_C E_C^{NS} = u_C^0 (1 - \varphi_{NS}) E_C^{NS} \geq u_C^0 E_C^0$, so the clean return weakly rises. Since $1 - \varphi_{NS} \leq 1$, it also gives $E_C^{NS} \geq E_C^0$, so the dirty return weakly falls. The strict pooled inequality $u_C^0 E_C^0 - c_C > \max\{0, u_D(I - E_C^0) - c_D\}$ therefore carries over. □

L.12 Support, verification and funding

The subsidy thresholds determine the incentive payment; the funding account must also make it available. Because $h_*'(s_*) > 0$, a finite additive clean payment is sufficient for fixed

payoff bounds, fixed competing payments and finite costs. Its date-t value and date-$t+1$ settlement remain as specified above. Zero clean support alone is a narrower case than fully private research if competing projects still receive payments.

The economic connection is between credible origin-based access, owned production opportunities and the research they reward. Installation and access contracts can transmit those opportunities to investors and later adopters. Research support supplements this return when the specified allocation calls for it; generic public research funding does not itself depend on certification.

Zero mineral-research grants still leave the cost of maintaining credible access. Verification services new physical deliveries, resolves claims on earlier deliveries and retains origin evidence permanently. With newly handled mass Q_t, permanent records D_t and a bounded number n_t of new records per physical unit, its conservative workload is

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

All three workload coefficients are positive. Paid verification research improves processing and storage productivity Z_t. Starting from installed and paid productivity $Z_0 > 0$, if throughput grows no faster than G^t, with $G > 1$, then productivity growth $g_V \geq G$ keeps W_t/Z_t bounded, including the permanent record stock. This supports a bounded service-capital requirement while retaining positive costs for scientists, laboratories, deployment, hardware, replacement and maintenance. The result relies on productivity improvements applying to old-record storage. Financial claims expire according to their contracts; the evidence remains.

Funding separates these resource costs from research incentives. Procurement escrow pays the supplier. Separate origin-and-vintage trusts secure liabilities until the claims close, and unresolved collateral cannot fund another cohort. Initial infrastructure and inherited claims, new deposits, realized losses and the cost of tied capital all have explicit funding. Fees pay part of the system; a stated outside endowment and unconditional appropriation finance the balance, including any research success payments. Paid adoption, local services and household support have their own entries. The construction establishes feasible funding under the stated bounds. It does not imply that certification receipts alone finance every program, or estimate a national fiscal envelope (see Online Appendix O.1).

L.13 From the available standard to government action

Paid adoption connects the successively improved clean standard to the legal production and tax opportunity considered in the government's decision. Copying the disclosed design costs no royalty, but equipment, site adaptation, custody and authorized market access must be financed. A successful intervention at date t installs the standard disclosed at t and first produces legally at $t+1$. An already legal adopter producing at t uses the

preceding standard $L_{t-1} = L_t/g$; the initial standards at dates -1 and 0 are inherited technologies.

North–South demand affects both alternatives through these market prices and the incomes earned by the region's legal and illegal operators. In the faithful CES specialization, the certified expenditure floor and dirty expenditure ceiling describe the two industry opportunities; in the contract markets, the actual regional and VCG values supply them. Legal tax receipts must fund actual oversight and the announced household livelihood contract before leaving a net fiscal return. Advance reserves pay for the interval before taxes arrive. Section M studies when this funded legal opportunity makes intervention preferable to protection income and waiting, allowing failure and another attempt. The same anticipated market path determines private innovation values, legal and illegal incomes, and government continuation values; reform retains the continuing obligations of the original funded parties.

M Government action and the transition

The industry changes developed in Section L alter a government's return from controlling an illegal mining region. Tolerating the illegal operation provides protection income. Establishing control costs resources and can fail, but gives a legal operator access to paid production opportunities and creates a tax base. The government compares the legal taxes left after local services and household support with the income it gives up by acting. I derive these two flows together, then solve the intervention comparison with waiting and repeated attempts. The resulting choices feed back into world prices, production and the prospects for conservation.

The North and South in Section L identify buyers; the government here belongs to the producing jurisdiction. A legal installation, subject to its authorized market access, can serve Northern demand for certified supply and Southern buyers who choose it. Illegal output reaches Southern outside demand and any fraudulent route actually available; its owner's income follows the actual selling contract. These market opportunities enter the local comparison through the net income each operator receives.

M.1 Tax revenue and protection income at the same prices

I consider a type-a government in a region of scale $w_r > 0$. One region has measure zero in the world economy and takes the anticipated public market path as given. Its legal and illegal alternatives are valued at the same mineral and final-product prices. Establishing control permits the legal operator to install the disclosed standard under Section L.13's paid contracts. An intervention at date t can first yield legal production at $t+1$. The legal opportunity therefore includes the installation bill and the adoption lag.

Let $\pi^{C}_{a,t}$ be the legal operator's optimized operating surplus, after its production costs and contractual obligations. Let $\pi^{D,\mathrm{disp}}_{a,t}$ denote the illegal actor's distributable net receipts, including net fraud income only when that actor owns it. The government's protection or bribe payment comes from these receipts. These are monetary flows per period in the outside-good numeraire. They use the actual regional production and selling problems. For example, if a region operates $n_{a,t}$ copies of Section L.6's explicit VCG cell, their optimized operating surplus is

$$\pi^{\mathrm{VCG}}_{a,t} = n_{a,t} \max_{q \in \mathcal{Q}_V} \left\{ W^V(q; p_t) - \frac{k_V q^2}{2} \right\} = n_{a,t} \left(p_t R_V - \frac{k_V}{2} \right) \tag{234}$$

Here $\mathcal{Q}_V$ is the cell's feasible inventory set, W^V its complete selling continuation including genuine retention, and the last equality holds on the stated strict branch where $q_V = 1$. Thus its profit-tax contribution is $\xi_a n_{a,t}(p_t R_V - k_V/2)$. The tax base is the optimized surplus; neither gross receipts nor final output can replace it. Where a jurisdiction operates both retail and bulk activities, their actual fiscal contributions must be added.

The government receives a profit-tax share $\xi_a \in (0,1)$ of the legal operating surplus. Let $F^S_{a,t} \geq 0$ be its local-service payments and $F^L_{a,t} \geq 0$ its livelihood-support payments. The two public income flows are

$$\begin{aligned} R_{a,t} &= \xi_a \pi^C_{a,t} - F^S_{a,t} - F^L_{a,t} \\ b_{a,t} &= \text{actual protection payment from } \pi^{D,\text{disp}}_{a,t} \end{aligned} \tag{235}$$

A fixed sharing contract, for example, would give $b_{a,t} = \zeta_a \pi^{D,\text{disp}}_{a,t}$, with $0 \leq \zeta_a < 1$. The general comparison below uses the actual payment $b_{a,t}$ and does not impose this particular sharing rule. It also keeps the admission organizer's income with its owner. Public adoption creates the local operator's income and tax base.

A proportional profit tax preserves the legal operator's static production choice: it multiplies its operating objective by $1 - \xi_a > 0$. The tax rate is a maintained fiscal rule in this model. The government's choice concerns intervention and its timing. This productive basis for fiscal incentives connects the mechanism to weak states and legal and fiscal capacity [Acemoglu, 2005, Besley and Persson, 2009].

Let c_t be the common income scale: $c_t = 1$ in the fixed-expenditure benchmark and $c_t = L_t$ in the connected growing-income classes. Dividing both alternatives by the same $w_r c_t$ gives

$$\widehat{R}_{a,t} = \frac{R_{a,t}}{w_r c_t}, \qquad \widehat{b}_{a,t} = \frac{b_{a,t}}{w_r c_t}, \qquad \widehat{\pi}^C_{a,t} = \frac{\pi^C_{a,t}}{w_r c_t} \tag{236}$$

The normalized service and support payments are $\widehat{F}^S_{a,t} = F^S_{a,t}/(w_r c_t)$ and $\widehat{F}^L_{a,t} = F^L_{a,t}/(w_r c_t)$. Suppose optimized legal surplus has a floor $\widehat{\pi}^C_{a,t} \geq \underline{\pi}_a > 0$, and these payments are bounded by $f_{S,a}$ and $f_{L,a}$. I impose the funding condition $0 < f_{S,a} + f_{L,a} < \xi_a \underline{\pi}_a$. Then

$$\widehat{R}_{a,t} \geq \underline{R}_a := \xi_a \underline{\pi}_a - f_{S,a} - f_{L,a} > 0 \tag{237}$$

The fiscal opportunity is therefore measured after funding the services and household contract needed to implement legal production.

Cash must also be available before taxes arrive. The normalized cost $k_a > 0$ of one specified intervention includes force, preparation and installation, advance-finance costs, and temporary household support where specified. An action-independent outside appropriation or initially endowed facility supplies liquidity, with a separate reserve for one local settlement period. Its common income cancels when the government compares acting with waiting; spending k_a still sacrifices numeraire value. Initial adoption is charged once. Continuing replacement, services and support during failed interventions retain their own funded entries.

These local accounts use the shared fees, outside-good endowment, initial reserves and unconditional appropriations in Section L.12. Earlier certificate claims remain obligations of their original signers and funded vintage trusts after reform. The new government does not acquire those debts through a change of control.

M.2 Solving the intervention comparison

The government of an ungoverned region chooses before the current protection payment. Waiting yields $\widehat{b}_{a,t}$ and preserves the option to act later. Acting costs k_a, forgoes that payment and succeeds with probability $P_a \in (0,1)$. Success opens legal production at the next date; failure leaves the region ungoverned, with the option to try again. Current physical production is unchanged by the attempt. The intervention technology and its success probability are given; the decision to use it is endogenous.

Let $U_{a,t}$ and $V_{a,t}$ be the government's discounted values, normalized by $w_r c_t$, when its region is governed and ungoverned. If its discount factor in monetary units is $\delta_G \in (0,1)$, the normalized continuation factor is $\widehat{\delta}_G = \delta_G$ at fixed expenditure and $\widehat{\delta}_G = \delta_G g < 1$ when $L_{t+1} = gL_t$. On the anticipated actual public path,

$$\begin{aligned} U_{a,t} &= \widehat{R}_{a,t} + \widehat{\delta}_G U_{a,t+1} \\ V_{a,t} &= \max\Big\{\widehat{b}_{a,t} + \widehat{\delta}_G V_{a,t+1}, \\ &\qquad\qquad -k_a + \widehat{\delta}_G[P_a U_{a,t+1} + (1-P_a)V_{a,t+1}]\Big\} \end{aligned} \tag{238}$$

The gain from acting rather than waiting is

$$\Delta_{a,t} = -k_a - \widehat{b}_{a,t} + \widehat{\delta}_G P_a (U_{a,t+1} - V_{a,t+1}) \tag{239}$$

The first two terms are the present sacrifice: the attempt costs resources and gives up current protection income. The last term values the chance of establishing legal control next period. The value of remaining ungoverned includes tomorrow's protection income and the option to intervene later.

M.2.1 A constant-flow benchmark

I first hold the normalized flows and intervention terms constant. For this calculation write R, b, k, P for their type-specific values. The legal value, the value of waiting forever, and the value of attempting intervention after every failure are respectively

$$U = \frac{R}{1-\widehat{\delta}_G}, \qquad V^{\text{wait}} = \frac{b}{1-\widehat{\delta}_G}, \qquad V^{\text{try}} = \frac{-k + \widehat{\delta}_G P U}{1 - \widehat{\delta}_G + \widehat{\delta}_G P} \tag{240}$$

The last expression solves $V^{\text{try}} = -k + \widehat{\delta}_G[PU + (1-P)V^{\text{try}}]$. After failure, the government returns to the same choice with continuation factor $\widehat{\delta}_G(1-P)$, leaving $1 - \widehat{\delta}_G + \widehat{\delta}_G P$ in the denominator. The numerator charges the attempt now and discounts the value of legal income beginning next period if the attempt succeeds. Since waiting and failure return to the same state in this benchmark, the optimal ungoverned value is $\max\{V^{\text{wait}}, V^{\text{try}}\}$.

Substituting the two income streams and collecting terms gives the threshold together with its opportunity cost:

$$\begin{aligned} V^{\text{try}} \geq V^{\text{wait}} \quad &\iff \quad \frac{-k + \widehat{\delta}_G P R/(1-\widehat{\delta}_G)}{1 - \widehat{\delta}_G + \widehat{\delta}_G P} \geq \frac{b}{1-\widehat{\delta}_G} \\ &\iff \quad \widehat{\delta}_G P(R-b) \geq (1-\widehat{\delta}_G)(k+b) \end{aligned} \tag{241}$$

Intervention is strictly preferred under a strict inequality; at equality the government is indifferent. Let f denote the constant normalized service and livelihood payments together, so $R = \xi_a \widehat{\pi}^C - f$. Solving for the legal operating surplus yields

$$\boxed{\widehat{\pi}^C \geq \widehat{\pi}^{C,*}(b) := \frac{1}{\xi_a}\left[f + b + \frac{1-\widehat{\delta}_G}{\widehat{\delta}_G P}(k+b)\right]} \tag{242}$$

The tax base must cover services and livelihoods, replace protection income, and compensate the government for costly, delayed and uncertain control. Under the fixed-share example, substituting $b = \zeta_a \widehat{\pi}^{D,\text{disp}}$ makes the comparison directly between optimized legal and illegal operating opportunities. Higher intervention costs or protection income raise the required legal surplus; a higher success probability lowers it. Prices enter both alternatives through their actual operating payoffs.

M.2.2 The threshold on the changing market path

As research and governments change production, the public path evolves. The date-t choice must then use $\Delta_{a,t}$ in (239). A fixed-flow annuity cannot replace its continuation values. The same fiscal calculation nevertheless gives a sufficient intervention threshold for the evolving economy.

Write $D_{a,t} = U_{a,t} - V_{a,t}$. Subtracting the two Bellman equations gives

$$\begin{aligned} D_{a,t} = \min\{&\widehat{R}_{a,t} - \widehat{b}_{a,t} + \widehat{\delta}_G D_{a,t+1} \\ &\widehat{R}_{a,t} + k_a + \widehat{\delta}_G(1-P_a) D_{a,t+1}\} \end{aligned} \tag{243}$$

Suppose that from a date T onward, legal revenue has the funded floor $\underline{R}_a$ and normalized protection income is at most b_*. Comparison with the constant-flow problem then gives

$$D_{a,t} \geq \min\left\{\frac{\underline{R}_a - b_*}{1-\widehat{\delta}_G}, \quad \frac{\underline{R}_a + k_a}{1-\widehat{\delta}_G + \widehat{\delta}_G P_a}\right\}, \qquad t \geq T \tag{244}$$

Inserting either branch into the intervention advantage gives the same strict fiscal test. Solving it for the maximum protection flow gives

$$\boxed{b_* < b_a^{\text{crit}} := \frac{\widehat{\delta}_G P_a \underline{R}_a - (1-\widehat{\delta}_G) k_a}{1 - \widehat{\delta}_G + \widehat{\delta}_G P_a}} \tag{245}$$

Every government of a still-ungoverned region of that type then strictly prefers intervention at every date from T. This is a sufficient bound on the changing path; its exact first intervention date is determined by the Bellman equation.

The critical protection flow is positive precisely when

$$k_a(1-\widehat{\delta}_G) < \widehat{\delta}_G P_a \underline{R}_a = \widehat{\delta}_G P_a[\xi_a \pi_a - f_{S,a} - f_{L,a}] \tag{246}$$

If $0 \leq \widehat{b}_{a,t} \to 0$ uniformly over feasible population paths, a finite T satisfies (245). The bound retains future reform in the value of waiting and permits the two income streams to fluctuate before it is reached. Theorem M.1 collects the two thresholds.

Theorem M.1 (The intervention threshold)**.** *A government intervenes when the legal tax base, net of services and livelihood support, exceeds protection income by enough to pay for costly, delayed and uncertain control.*

1. *With constant normalized flows, attempting intervention after every failure beats waiting forever exactly when $\widehat{\delta}_G P(R-b) \geq (1-\widehat{\delta}_G)(k+b)$, as in (241); equivalently, when the legal operating surplus meets the tax-base threshold (242).*

2. *On a changing path, suppose that from date T the funded legal flow is at least $\underline{R}_a$ and normalized protection income is at most b_*. If $b_* < b_a^{\text{crit}}$ in (245), every government of a still-ungoverned region of type a strictly prefers intervention at every date from T. The critical flow is positive exactly when (246) holds.*

(see Appendix Q.6)

Scope. Part 2 is a sufficient bound; the exact first intervention date follows from the Bellman equation (238). If $0 \leq \widehat{b}_{a,t} \to 0$ uniformly, a finite T exists. This holds in the growing-income constructions, where legal opportunity grows with L_t while protection income stays bounded. At fixed expenditure the protection-flow bound must be established separately, and outside the strict fiscal conditions waiting can remain optimal.

M.3 Market feedback and the transition

The growing legal standard of Sections L.10 and L.13 can deliver the relative fiscal opportunity used in the threshold. In the connected growing-income constructions, an unreformed region retains fixed dirty capacity, while prices and its actual fraud opportunities remain bounded on the stage domain. Its absolute protection income is therefore bounded. The legal adopter installs the growing standard with its paid lag, allowing a funded net fiscal return to remain a positive fraction of L_t. Protection income divided by L_t tends uniformly to zero. Condition (246) can thus hold on an open set with positive intervention, service and support costs.

This comparison allows scarcity to raise the returns of illegal producers who remain after other regions reform. Legal opportunity can still grow relative to their bounded protection

income. Additional demand conditions can make absolute dirty retail rents fall on the stated branch; that stronger result does not automatically sign a region's bulk activities. The physical extraction response also uses the cell's mineral input q_D, rather than its transformed final output Z_D (see Online Appendix, Proposition N.22 and Corollary N.23). At fixed expenditure the protection-flow bound must be established separately. A positive dirty expenditure share does not establish vanishing protection income, and outside the strict fiscal conditions waiting can remain optimal indefinitely.

Government choices then determine the transition of regions. Let $M^D_{a,t}$ and $M^C_{a,t}$ be the scale-weighted geographic masses of ungoverned and governed regions, distinct from research-owner masses and buyer counts. If fraction $u_{a,t}$ of the ungoverned mass intervenes,

$$\begin{aligned} M^D_{a,t+1} &= (1 - u_{a,t}P_a)M^D_{a,t} \\ M^C_{a,t+1} &= M^C_{a,t} + u_{a,t}P_a M^D_{a,t} \end{aligned} \tag{247}$$

This deterministic population law uses independent regional outcomes and exact aggregation in an atomless population. A common intervention shock would require its distribution to enter the state. Once attempts occur every date, remaining ungoverned mass declines by $1 - P_a$ per date, and each individual region succeeds in finite time almost surely.

One atomistic government takes the market path as given, while the actions of a positive population must generate that path. Successful control replaces the relevant ungoverned activity from the next date. Markets clear again, research owners' innovation values are repriced, and government continuation values are recomputed. Governments therefore affect one another through common mineral and final-product prices. This differs from the municipal network studied by Acemoglu et al. [2015].

The connected-equilibrium result (Theorem N.21) keeps these choices consistent with the stage market, research and finance. It requires the verified stage domain, funded verification and budget conditions, the strict research comparison using valid bounds on complete innovation values, and continuous bounded normalized flows and vintage obligations. In the growing-income classes, $\max(\delta_G, \delta_R)g < 1$ bounds discounted continuations. Under those conditions an infinite-horizon equilibrium exists, with waiting, intervention and mixing at indifference. Government availability permits the optimizing choice described above. The existence result does not establish global uniqueness (see Online Appendix O.2).

M.4 Funded livelihoods and protected sites

Political control changes the production opportunities available to households. To connect that change to conservation, each market outcome is returned to a feasible regional allocation through the implementation correspondence. Let $y_{rh}(q_{r,t})$ be a household's feasible net livelihood on the selected legal branch, after effort and its baseline alternatives are valued in the stated numeraire convention. A promised livelihood floor $\underline{y}_{rh}$ is implemented by

$$T_{rh,t} = [\underline{y}_{rh} - y_{rh}(q_{r,t})]_+ \tag{248}$$

The payment is fixed from the committed allocation and is independent of the household's subsequent marginal output choice. It is charged to local taxes or the unconditional outside appropriation, subject to the per-scale support bound used in the fiscal floor. Services, initial paid adoption, continuing replacement and support during failed interventions remain funded, including the advance reserve before legal taxes arrive. This makes the legal opportunity a feasible household allocation as well as a prospective government receipt.

The physical target in Assumption 1.6 is the safe allocation at designated protected sites under the asserted date or limiting criterion. The conservation construction requires the legal implementation set to enforce zero extraction disturbance in those sites. Paid legal installations use separately specified noncritical locations, and fixed research cores also operate outside the critical habitat. A certificate or a higher legal tax base does not supply this physical restriction.

Let an ungoverned region's protected-site disturbance be bounded by $w_r \bar{d}_a$. After intervention is strictly chosen by all remaining governments, the geometric decline of $M^D_{a,t}$ makes expected aggregate disturbance tend to zero; with the stated exact continuum aggregation, the same conclusion holds for the exact aggregate. Every strictly positive tolerance is reached after a finite date. For $P_a < 1$, positive ungoverned mass can nevertheless remain at each finite date. Almost-surely finite success of each individual region does not give a common finite success date for the atomless population.

Delay matters separately when damage accumulates.

Proposition M.2 (Transition and cumulative disturbance)**.** *Once every remaining government intervenes, the ungoverned mass falls geometrically and cumulative damage to protected sites is bounded. If a fraction $u_{a,t}$ of ungoverned mass intervenes, the masses follow (247); with attempts at every date the ungoverned mass falls by the factor $1 - P_a$ per date, and expected aggregate disturbance tends to zero. If all remaining types begin intervention by a common date T, cumulative disturbance stays within a specified allowance $\mathcal{B}$ whenever*

$$\sum_a \bar{d}_a M^D_{a,0} \left(T + \frac{1}{P_a} \right) \leq \mathcal{B} \tag{249}$$

(see Appendix Q.7)

Scope. Legal production must use the specified safe implementation set. For $P_a < 1$ positive ungoverned mass remains at every finite date, and fixed dirty cores outside protected habitat may remain productive.

The bound has the same expected or exact-continuum interpretation as the aggregate result. The term T covers waiting before intervention becomes universal, and $1/P_a$ covers disturbance during repeated attempts. This cumulative requirement is separate from convergence to zero and from finite crossing of a positive tolerance. Later reform cannot reverse damage treated as irreversible.

The route to conservation therefore combines a fiscal incentive to establish control with paid adoption, feasible household livelihoods and a physically safe location. Its timing

must meet the stated physical objective. The aggregate limit and transition-budget result identify what the mechanism establishes; they do not identify candidate extraction with forest loss or supply exact safety at every site at a common finite date.

N Supporting proofs and technical definitions

N.1 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.

N.1.1 Measure, scale and the cost infimum

Assumption N.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

$$s = 1 + 1/\eta, \qquad \alpha_m = 1/s = \eta/(1+\eta), \qquad \beta_r = A_r \kappa_r^{-\eta}, \qquad 0 < B = \int \beta_r \, d\mu^{\mathrm{reg}} < \infty. \tag{250}$$

All $q \geq 0$ are feasible, zero-output exit has cost and payoff zero, and each fiber attains the private implementation cost

$$\mathcal{C}_r(q) = \alpha_m \beta_r^{-1/\eta} q^s. \tag{251}$$

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

$$\overline{\mathcal{C}}(Q) = \inf\left\{ \int \mathcal{C}_r(q_r) \, d\mu^{\mathrm{reg}} : q_r \geq 0 \text{ measurable}, \quad \int q_r \, d\mu^{\mathrm{reg}} = Q \right\}. \tag{252}$$

The integral is an extended nonnegative integral, so the definition permits candidate allocations with infinite cost.

Theorem N.2 (Exact cost, supply and operator-profit aggregation)**.** *Under Assumption N.1, the cost infimum is attained and equals*

$$\overline{\mathcal{C}}(Q) = \alpha_m B^{-1/\eta} Q^s. \tag{253}$$

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

$$\begin{aligned} q_r^* &= \beta_r [r_0]_+^\eta, & \Pi_r^* &= \frac{\beta_r [r_0]_+^{1+\eta}}{1+\eta}, \\ Q^* &= B[r_0]_+^\eta, & \Pi^* &= \frac{B[r_0]_+^{1+\eta}}{1+\eta}. \end{aligned} \tag{254}$$

The aggregate maximum of $r_0 Q - \overline{\mathcal{C}}(Q)$ gives exactly these quantities and profits. The regional output, implementation-cost and profit functions are integrable.

Proof. Set $d\nu = \beta_r \, d\mu^{\mathrm{reg}}/B$, 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

$$\int \mathcal{C}_r(q_r) \, d\mu^{\mathrm{reg}} = \alpha_m B \int f_r^s \, d\nu \geq \alpha_m B (Q/B)^s. \tag{255}$$

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 $q \to \infty$. 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. □

Corollary N.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*

$$\sup_{(q_r)} \left\{ W\left(\int q_r \, d\mu^{\mathrm{reg}}; \Omega \right) - \int \mathcal{C}_r(q_r) \, d\mu^{\mathrm{reg}} \right\} = \sup_{Q \geq 0} \{ W(Q;\Omega) - \overline{\mathcal{C}}(Q) \}. \tag{256}$$

If W is finite valued and upper semicontinuous and $W(Q;\Omega) - \overline{\mathcal{C}}(Q)$ tends to minus infinity as $Q \to \infty$, 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 N.6, a

household reconstruction. This statement does not aggregate independent local nonlinear continuations into a representative monopoly.

Proof. For any regional allocation, Theorem N.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 (256). 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. □

The worked example is in Online Appendix O.5.

N.1.2 Contracts, accounting moments and invariance

Proposition N.4 (Distinct supply and profit moments)**.** *Consider finitely many groups satisfying Assumption N.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*

$$Q_j = p_j^{\eta} S_j, \quad \Pi_j = \frac{p_j^{1+\eta}}{1+\eta} H_j, \qquad S_j = \sum_{g \in j} B_g \omega_g^{\eta}, \quad H_j = \sum_{g \in j} B_g \omega_g^{1+\eta}. \tag{257}$$

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 (254) group by group and sum. The cost expression follows from the regional cost at the optimum. Strictness follows from the equality condition in (255); 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. □

The worked example is in Online Appendix O.6.

With nonnegative fees $p_j - r_g$, the identity

$$p_j Q_j = \Pi_j + \int \mathcal{C}_r(q_r)\, d\mu^{\mathrm{reg}} + \sum_{g \in j} (p_j - r_g) Q_g \tag{258}$$

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 $Q = \int \beta_r [r(r)]_+^{\eta}\, d\mu^{\mathrm{reg}}$ and $\Pi = \int \beta_r [r(r)]_+^{1+\eta}\, d\mu^{\mathrm{reg}}/(1+\eta)$ 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 $\overline{\mathcal{C}}_{\lambda B}(\lambda Q) = \lambda\overline{\mathcal{C}}_B(Q)$. If mineral is measured in units with $q' = uq$ and currency is measured with money$' = v$ money, then

$$p' = \frac{v}{u}p, \qquad \beta' = u^{1+\eta}v^{-\eta}\beta, \qquad \mathcal{C}'(q') = v\mathcal{C}(q), \qquad \Pi' = v\Pi. \tag{259}$$

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.

N.1.3 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 N.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: $\{r : F(r) \cap O \neq \varnothing\}$ 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 N.6 (Attained measurable household reconstruction)**.** *Under Assumption N.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 N.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 $\{r : F(r) \cap B(y_k, a) \neq \varnothing\}$. 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

$$d_k(r) < 2^{-n}, \qquad \|y_k - z_{n-1}(r)\| < 3\,2^{-n}. \tag{260}$$

Each such first-index choice is measurable, because its countable candidate sets are measurable. The recursion is feasible: since $\operatorname{dist}(z_{n-1}, F) < 2^{-(n-1)}$, 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 (260) 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. □

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 $\int c_r^g g_r \, d\mu^{\mathrm{reg}}$ does not imply finite physical input $\int g_r \, d\mu^{\mathrm{reg}}$. 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 N.7 (A fully constructive household branch)**.** *Let Assumption N.1 hold and set*

$$a_r = \left(\frac{c_r^g}{\kappa_r \alpha_m^2}\right)^{\alpha_m} A_r^{1-\alpha_m}, \qquad q = a_r g^{\alpha_m}. \tag{261}$$

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

$$\log(C_{hr} - c_{hr}) + \nu_{hr} \log(b_{hr} + b_{hr}^0), \qquad b_{hr} \geq 0. \tag{262}$$

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, $\int \sum_h (w_{hr} + c_{hr}) \, d\mu^{\mathrm{reg}} < \infty$, and $\int (\beta_r / c_r^g + c_r^g) \, d\mu^{\mathrm{reg}} < \infty$. 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:

$$\begin{aligned} \mathcal{C}_r &= \alpha_m \beta_r t^s, & p_r^g &= c_r^g, & \theta_r &= \alpha_m t^{1/\eta}, \\ g_r &= \frac{\alpha_m^2 \beta_r}{c_r^g} t^s, & H_r &= (1-\alpha_m)\mathcal{C}_r, & M_{hr} &= w_{hr} + \xi_{hr} H_r - \bar{p}_{hr} c_{hr}, \\ b_{hr} &= \left[\frac{\nu_{hr} M_{hr} - b_{hr}^0}{1+\nu_{hr}} \right]_+, & C_{hr} &= c_{hr} + \frac{M_{hr} - b_{hr}}{\bar{p}_{hr}}. \end{aligned} \tag{263}$$

At $Q = 0$ take the zero procurement quote or the separately feasible exit with the same surviving outside allocation. Procurement accounting obeys

$$\mathcal{C}_r = \underbrace{c_r^g g_r}_{\text{real productive input}} + \underbrace{H_r}_{\text{household operating income}}, \qquad \bar{p}_{hr} C_{hr} + b_{hr} = w_{hr} + \xi_{hr} H_r. \tag{264}$$

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

$$\theta q - (p_r^g - c_r^g) g = c_r^g g + \frac{1-\alpha_m}{\alpha_m} p_r^g g. \tag{265}$$

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 (261) 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 (263). 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 $0 \le b_{hr} \le M_{hr}$ and $0 < C_{hr} - c_{hr} \le M_{hr}/\bar{p}_{hr}$, so the assumed wealth, subsistence and price bounds give integrability. Consumption surplus is bounded below by $\epsilon/[p_{\max}(1+\overline{\nu})]$ and $b + b^0$ is bounded away from zero. Their logarithms thus have bounded negative parts; $\log^+ y \le y$ 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. □

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 (265).

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.

The worked example is in Online Appendix O.7.

N.1.4 Entry, finite coverage and subsistence restrictions

Proposition N.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*

$$V_r = \frac{\beta_r [r_0]_+^{1+\eta}}{1+\eta} - f_r. \tag{266}$$

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

$$Q = [r_0]_+^{\eta} \int \beta_r e_r \, d\mu^{\mathrm{reg}}, \qquad \Pi = \int [V_r]_+ \, d\mu^{\mathrm{reg}}, \tag{267}$$

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 $\int \beta_r \, d\mu^{\mathrm{reg}} < \infty$ and $\int f_r \, d\mu^{\mathrm{reg}} < \infty$.

If $\beta_r = \beta_0 > 0$ and $f_r = f > 0$ throughout an atomless population of mass m, define

$$r_E = \left(\frac{(1+\eta) f}{\beta_0} \right)^{1/(1+\eta)}, \qquad q_E = \beta_0 r_E^{\eta}. \tag{268}$$

The attained aggregate fixed-output cost is

$$\overline{\mathcal{C}}^{\,f}(Q) = \begin{cases} r_E Q, & 0 \le Q \le m q_E, \\ m f + \alpha_m (m\beta_0)^{-1/\eta} Q^s, & Q \ge m q_E. \end{cases} \tag{269}$$

For $0 < Q < mq_E$, exactly mass Q/q_E operates at output q_E per active region. At and above mq_E, all regions operate with output Q/m. At receipt r_E, any aggregate output in $[0, mq_E]$ 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

$$F(n) = fn + \alpha_m \beta_0^{-1/\eta} Q^s n^{-1/\eta}, \qquad 0 < n \le m. \tag{270}$$

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 $Q > 0$, $F(n) \to \infty$ 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 mq_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. □

Proposition N.9 (Every heterogeneous entry total has an attained reconstruction)**.** *Retain the atomless population and power variable costs of Assumption N.1, and allow any measurable finite fixed costs $f_r \ge 0$. Define the entry thresholds and integrated operating profit by*

$$\tau_r = \left(\frac{(1+\eta) f_r}{\beta_r} \right)^{1/(1+\eta)}, \qquad P(p) = \int \left[\frac{\beta_r p^{1+\eta}}{1+\eta} - f_r \right]_+ d\mu^{\mathrm{reg}}, \quad p \ge 0. \tag{271}$$

For every finite $Q > 0$ there is a unique $p > 0$ with

$$p^\eta \int_{\{\tau_r < p\}} \beta_r \, d\mu^{\mathrm{reg}} \le Q \le p^\eta \int_{\{\tau_r \le p\}} \beta_r \, d\mu^{\mathrm{reg}}. \tag{272}$$

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

$$\overline{\mathcal{C}}^{\,f}(Q) = pQ - P(p) = \sup_{u \geq 0}\{uQ - P(u)\}. \tag{273}$$

Its output, variable cost and paid fixed costs are integrable, even if $\int f_r \, d\mu^{\mathrm{reg}} = \infty$. 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 $d\nu = \beta_r \, d\mu^{\mathrm{reg}}$. 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 $p \to \infty$, the sets $\{\tau < p\}$ increase to the full population up to null sets, so $Q^-(p) \to \infty$. 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 individually 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. □

The worked example is in Online Appendix O.8.

For completeness, a finite measurable coverage menu can be handled without extrapolating a power envelope. In branch $a \in \{1, \ldots, K\}$ 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

$$q_r^a = \min\{U_r^a, \max\{L_r^a, \beta_r^a [r_0]_+^\eta\}\}. \tag{274}$$

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.

The worked example is in Online Appendix O.9.

N.2 Common-price equilibrium, composition and four accounts

N.2.1 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 $\mathscr{W}_\Sigma(q - d \mid q, \Omega) = p_D(q - d)$ 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 participants after obtaining it, and optimize all feasible informed mitigation before the printed silence loss is used.

N.2.2 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

$$z < 1 - x - b_k[r_k + (1-r_k)e], \qquad z \sim U[0, c_k]. \tag{275}$$

The right side includes the dirty receipts forgone through false admission. Bayes consistency gives

$$\pi = \left[\frac{1 - x - b_k[r_k + (1-r_k)\min\{a_k\pi, 1\}]}{c_k}\right]_0^1. \tag{276}$$

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 (329), 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

$$O_k = hX_k \int_0^{c_k\pi_k} \frac{c_k\pi_k - z}{c_k}\, dz = \frac{hc_k\pi_k^2 X_k}{2}, \tag{277}$$

$$F_k^{\text{real}} = hX_k \int_0^{c_k\pi_k} \frac{z}{c_k}\, dz = \frac{hc_k\pi_k^2 X_k}{2}. \tag{278}$$

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.

N.2.3 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 $e_k < 1$,

$$R_{k,\pi} = r_k l_{Lk} - A_k e_k > 0 \tag{279}$$

except the harmless zero-loss limit. Every term of F_{cw} and H_{cw} is nonincreasing in x, and

$$\frac{d}{d\pi}\frac{1+\phi+R_k(\pi)}{1-\pi} = \frac{R_{k,\pi}(1-\pi)+1+\phi+R_k}{(1-\pi)^2} > 0. \tag{280}$$

Thus $T' \leq 0$. Put $Y_0(x) = B_D x^\eta - H_{\mathrm{cw}}(x) > 0$. Its derivative is $\eta B_D x^{\eta-1} - H'_{\mathrm{cw}} > 0$, and

$$J'(x) = \frac{Y_0 + xY'_0}{T} - \frac{xY_0T'}{T^2} > 0. \tag{281}$$

Theorem N.27 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.

The complete analytical construction is in Online Appendix O.13.

N.2.4 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

$$Q_{Ck} + Q_{Dk} = X_k(1-\pi_k p_k) + X_k \pi_k p_k + Y_k. \tag{282}$$

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 (336). 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

$$hX_k = h(1-\pi_k)X_k + p_D\pi_k X_k + hb_k p_k \pi_k X_k + F_k^{\text{real}} + O_k. \tag{283}$$

Indeed the leftover spread is $h\pi_k X_k(1-x-b_k p_k) = hc_k\pi_k^2 X_k$ 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

$$p_D\sum_k Y_k + p_D\sum_k \pi_k X_k = p_D\sum_k Q_{Dk}. \tag{284}$$

Subtracting their private implementation cost gives (335). The clean suppliers' receipts similarly are $h\sum_k Q_{Ck}$. The fee budget is separate:

$$h\phi X = \text{paid ordinary service and capital bills} + \text{retained funded service balance}. \tag{285}$$

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.

N.2.5 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.

N.3 Proofs for the active nonlinear world

N.3.1 The cell, its regional cost, and its true physical account

All markets below are atomless replications of cells. A continuum menu cell contains the unit-demand buyer population specified below; its receipts and resale allocation are derived explicitly. 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

$$\int \frac{k}{2} z(a)^2 \, da \;\geq\; \frac{k}{2n_r} Q_r^2, \quad \int z(a)\, da = Q_r, \tag{286}$$

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.

N.3.2 A continuum menu cell with active resale and 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

$$P(u) = \begin{cases} 1 - 8u/3, & 0 \le u \le 1/4, \\ 4(1-u)/9, & 1/4 \le u \le 1. \end{cases} \tag{287}$$

Write

$$a = \frac{6+\sqrt{6}}{48}, \quad b = \frac{1+\sqrt{6}}{8}, \quad s = 1 - \frac{16a}{3}, \quad c = \frac{8}{3}a^2, \qquad R(u) = uP(u). \tag{288}$$

The least concave majorant is $\overline{R}(u) = c + su$ 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

$$V(m) = \begin{cases} (M - \tau_0)m, & 0 \le m \le r, \\ (1-\rho)[\overline{R}(m-r) + r(M-\tau_0)] + \rho m P(m-r), & r < m \le 1 + r. \end{cases} \tag{289}$$

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 $(1-\rho)[\overline{R}(m-d) + (M-\tau_0)d] + \rho m P(m-r)$. Its slope in d is strictly positive because all slopes of $\overline{R}$ are at most one. Thus $d = r$; the two contact thresholds attain the bound and clear the same secondary market. For $m \le r$, 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. The separated buyer groups permit this explicit derivation. A volume-only resale formula for general regulated distributions would require a separate equilibrium argument.

Introduce

$$d = \frac{16\rho}{3}, \quad \alpha = (1-\rho)s + \rho(1 + 8r/3), \quad m_0 = 41/200, \quad k = \frac{\alpha - dm_0}{m_0}, \tag{290}$$

and the independent, genuine retained conversion

$$H_\delta(y) = \begin{cases} h_0 y - y^2/(2\delta), & 0 \le y \le \delta h_0, \\ \delta h_0^2/2, & y \ge \delta h_0, \end{cases} \qquad \delta = 1/1000, \quad h_0 = 1/5. \tag{291}$$

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 N.10 (Global continuum-menu solution on an explicit price interval)**.** *For every $p \in I = [99/100, 101/100]$, the global maximum in (323), with V as in (289), is unique and equals*

$$\begin{aligned} m(p) &= \frac{p\alpha - k\delta(h_0 - \alpha)}{k(1+\delta d) + pd}, \\ y(p) &= \delta[h_0 - \alpha + dm(p)], \quad q(p) = m(p) + y(p), \\ \pi(p) &= p[V(m(p)) + H_\delta(y(p))] - kq(p)^2/2. \end{aligned} \tag{292}$$

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 $a + r \le m \le 1/4 + r$ 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 $1/4 + r \le m \le b + r$, 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 $m \le r$ 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 $p\delta h_0^2/2 \le 2.02 \cdot 10^{-5}$, 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 (292); 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 $p(1-\rho)[\overline{R}(u) - uP(u)] > 0$. 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. □

Let $I_P(u) = \int_0^u P(v)\,dv$. Its explicit formula is $u - 4u^2/3$ for $u \le 1/4$, 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

$$Z_C(p) = rM + (1-\rho)\{I_P(a) + \lambda[I_P(b) - I_P(a)]\} + \rho I_P(u) + H_\delta(y(p)). \tag{293}$$

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_C by q, V or $\overline{R}$ would be wrong. Direct differentiation gives

$$\begin{aligned} m'(p) &= \frac{k(\alpha + \delta d h_0)}{[k(1+\delta d) + pd]^2} > 0, \\ q'(p) &= (1+\delta d)m'(p) > 0, \\ Z_C'(p) &= m'(p)\left[(1-\rho)\frac{I_P(b) - I_P(a)}{b-a} + \rho P(u) + \delta d H_\delta'(y)\right] > 0, \\ \pi'(p) &= V(m(p)) + H_\delta(y(p)) > 0. \end{aligned} \tag{294}$$

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.

N.3.3 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 N.11 (Nonlinear finite continuation and actual production)**.** *Put* $R_h = h + (1 - h)^3/3$. *The actual all-offered continuation is*

$$W^V(q;p) = \begin{cases} phq, & 0 \le q < 1, \\ p[h(q-1) + R_h], & 1 \le q < 2, \\ phq, & q \ge 2. \end{cases} \tag{295}$$

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 $1 \le q < 2$, 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. □

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 the continuum menu cell. Whole-package capacity also explains why general differentiable supply elasticities cannot simply be asserted for every finite auction.

N.3.4 The equilibrium operator, reconstruction and scale

Theorem N.12 (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*

$$p_- \sum_j n_j Z_j(p_-) < E_R < p_+ \sum_j n_j Z_j(p_+), \qquad [p_-, p_+] \subseteq I. \tag{296}$$

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

$$Q_j = n_j q_j(p), \qquad \Pi_j = n_j \pi_j(p), \qquad Z_R = \sum_j n_j Z_j(p). \tag{297}$$

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. □

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 the finite VCG cell and strictly increasing for the continuum menu cell. 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 $n_{j,t} = L_t \widehat{n}_{j,t}$, $E_{Rt} = L_t e_{Rt}$ inside this compact domain, normalized profits and physical controls are uniformly bounded. If $L_{t+1}/L_t \le g$ 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.

N.3.5 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 N.12. 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

$$\begin{aligned} R_C &= hK(x)X - E_C, \\ R_D &= hxY - E_D - \varepsilon Y/(Y + \gamma Z_R(p)), \\ R_R &= pZ_R(p) - E_R - \varepsilon\gamma Z_R(p)/(Y + \gamma Z_R(p)). \end{aligned} \tag{298}$$

Here X, Y are the *same* endogenous quantity functions as in the bulk certification theorem; Z_R is (293) or its finite VCG counterpart, summed over actual cell masses.

Proof of Theorem N.29. The utility in (339) 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 (298). 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. □

This proof is a nonempty extension, since the bulk class has strict interior examples, Lemmas N.10 and N.11 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 (298) by L gives the same normalized system; its regularity and uniform bounds persist while that state remains in the verified compact domain. If $Y/L \to 0$ 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.

N.3.6 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 $L\widehat{n}_j$, with a positive clean retail core. Write $\widehat{n}_j = n_j/L$ and $\widehat{n} = (\widehat{n}_j)_j$ for their normalized masses. Their normalized final output is $Z(p;\widehat{n}) = \sum_j \widehat{n}_j Z_j(p)$.

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 (339). 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

$$\begin{aligned} t_v(x) &= \frac{1+\phi+\mathcal{R}_v(x)}{1-\pi(x)}, \\ X &= h^\eta \left[Lb_s + \frac{\beta_v}{1-\pi(x)} \right], \\ Y &= h^\eta \left[B_D x^\eta - \frac{\beta_v \pi(x)}{1-\pi(x)} \right]. \end{aligned} \tag{299}$$

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 N.13 (Uniform mixed-scale connected operator)**.** *Choose a compact state rectangle $s = (b_s, B_D, \widehat{n}, e_c, e_d, e_r)$ 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*

$$\begin{aligned} 0 &= h^{1+\eta}[(1+\phi)b_s + \lambda\beta_v t_v(x)] - e_c, \\ 0 &= xhY - e_d - \frac{\varepsilon_u \lambda Y}{\lambda Y + \gamma Z(p;\widehat{n})}, \\ 0 &= pZ(p;\widehat{n}) - e_r - \frac{\varepsilon_u \lambda \gamma Z(p;\widehat{n})}{\lambda Y + \gamma Z(p;\widehat{n})} \end{aligned} \tag{300}$$

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 continuum-menu 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

$$\sup \|D\Phi_{\beta_v,\varepsilon_u} - D\Phi_{0,0}\| \le C(\beta_v + \varepsilon_u). \tag{301}$$

For the retail price alone, with $p^0(\widehat{n}, e_r)$ *the scalar retail root, the stronger bound* $\sup \|D(p - p^0)\| \le C\varepsilon_u$ *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 (299) 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 (300). Its denominator stays positive at $\lambda = 0$ because Z has a positive core.

At $\beta_v = \varepsilon_u = 0$ the explicit baseline is

$$h^0 = \left[\frac{e_c}{(1+\phi)b_s}\right]^{1/(1+\eta)}, \quad p_D^0 = \left[\frac{e_d}{B_D}\right]^{1/(1+\eta)}, \quad x^0 = p_D^0/h^0, \quad p^0 Z(p^0; \widehat{n}) = e_r.$$

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 $\|I - A(s)^{-1} D_u R_{0,0}(u; s, \lambda)\| \le 1/4$. Compactness and smoothness of (300) 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 $e_r = Z(1; \widehat{n})$. 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, inverse Jacobians and the price displacement bound imply (301) 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 be sufficient. □

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 N.4.

N.3.7 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 N.14 (Funded truthful retail admission)**.** *On the compact stage domain of Theorem N.13, 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 $\overline{q}$ bound cell extraction and $\overline{F}$ 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 $\overline{F}$. Extra ledger units without measured backing are inadmissible. The continuum menu cell has buyer mass $1 + r$, 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 $[\underline{W}, \overline{W}]$. For the pre-report, no-invoice contract here take $\underline{W} = 0$; priced residual liabilities would instead remain in that bound. Formula (302), and its same-stock false-origin counterpart, give

$$G(a) \le \overline{F} + \overline{W} - \underline{W} =: G_* < \infty.$$

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 is sufficient 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. □

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.

N.3.8 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 selection. At every compared history the stipulated continuation must have an equilibrium. Let Ω be the market state. The exact organizer's deviation comparison is

$$F_C(a) - k_f(a) + \mathscr{W}_\Sigma(q - d(a) \mid q, \Omega) - \mathscr{W}_\Sigma(q \mid q, \Omega), \tag{302}$$

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 N.11 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 (338) falls at every fixed price. Equations (294) then raise remaining continuum-menu 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.

N.4 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 expenditure. 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.

N.4.1 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

$$M_{a,t}^D = \int_{\mathcal{R}_a} w_r 1_{\{r \text{ ungoverned at } t\}}, d\mu_a^{\text{reg}}, \qquad M_{a,t}^C = M_a - M_{a,t}^D, \qquad 0 < \sum_a M_a = \epsilon < \bar{\epsilon}.$$

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

$$\begin{aligned} B_t^C &= L_t \left(c_0 + \sum_a \frac{\omega_a^C}{g} M_{a,t}^C \right), & B_t^D &= d_0 + \sum_a \omega_a^D M_{a,t}^D, \\ \widehat{B}_t^C &= B_t^C / L_t, & \widehat{B}_t^D &= B_t^D / L_t. \end{aligned} \tag{303}$$

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 (303) denotes the mass of *complete independent installations*; its physical cost coefficient is obtained from the cell's stated normalization. In this retail construction $\widehat{n}_{C,t} = \widehat{B}_t^C$ and $\widehat{n}_{D,t} = \widehat{B}_t^D$ 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 replace whole VCG buyers with a continuum of unit-demand buyers.

Assumption N.15 (Verified stage interface for the dynamic construction)**.** 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 $Z_j(p)$:

$$p \sum_j \widehat{n}_{j,t} Z_j(p) = e_t. \tag{304}$$

The left side is strictly increasing on a common price interval $[\underline{p}, \overline{p}]$, 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 (303), including $\widehat{B}^D = 0$. 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 N.16 (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 N.13: with* $x = p_D/h$,

$$\begin{aligned} 0 &= h^{1+\eta}[(1+\phi)b_s + \lambda\beta_v t_v(x)] - e_c, \\ 0 &= xhY - e_d - \chi\frac{\lambda Y}{\lambda Y + \gamma Z}, \\ 0 &= pZ - e_r - \chi\frac{\lambda\gamma Z}{\lambda Y + \gamma Z}, \end{aligned} \tag{305}$$

where $Z = \sum_j \widehat{n}_j Z_j(p)$, $t_v = (1+\phi+R_v)/(1-\pi)$, *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 N.13 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 (303). 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. □

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.

N.4.2 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

$$D_{t+1} = D_t + n_t Q_t, \qquad 0 \leq n_t \leq n. \tag{306}$$

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

$$W_t = a_0 Q_t + a_{\text{hist}} \sum_{s=1}^{T_H} Q_{t-s} + a_D D_t, \qquad a_0, a_{\text{hist}}, a_D > 0. \tag{307}$$

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

$$Z_{t+1} = g_V Z_t, \qquad g_V = 1 + \gamma_V\left(1 - e^{-\nu_V S_V}\right) > 1. \tag{308}$$

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 N.17 (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 $D_t \le D_0 + n\bar{Q}(G^t - 1)/(G - 1)$, plus the specified finite initial cohorts if needed. Thus, for a finite constant C_{init},

$$\frac{W_t}{Z_t} \le \frac{\bar{Q}}{Z_0}\left(a_0 + a_{\text{hist}}\sum_{s=1}^{T_H} G^{-s} + \frac{a_D n}{G-1}\right)\left(\frac{G}{g_V}\right)^t + \frac{C_{\text{init}}}{Z_0 g_V^t} =: \widehat{W}_t \le \bar{W}.$$

Install $K \ge \bar{W}/\upsilon$ units of hardware, where $\upsilon > 0$ 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 $\gamma_V \ge (G-1)/(1 - e^{-\nu_V S_V})$ is finite. This proves feasibility, not that these bills pay themselves from a zero fee. □

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 $g_V > 1$ is sufficient 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 N.16. An invoice cohort has value $I_t \le \bar{I}c_t$. 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 N.18 (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 $b\sum_{s=0}^{T_H-1} I_{t-s} \le T_H b\bar{I}c_t$. 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. □

N.4.3 Actual research prizes, including surviving private liabilities

Let $\mathcal{O}^s_{i,t+1}, \mathcal{O}^f_{i,t+1}$ 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^z_i be its dated installation/access bill, B^z_i its new collateral deposit, and $C^z_{i,t+u}$ 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

$$J_{i,t} = \delta_R \, \mathbb{E}[\Delta \mathcal{O}_{i,t+1} - \Delta K_i - \Delta B_i] - \sum_{u=2}^{T_H+1} \delta_R^u \mathbb{E} \Delta C_{i,t+u} + \delta_R^{T_H+1} \mathbb{E} \Delta V^{\text{other}}_{i,t+T_H+1}. \tag{309}$$

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 N.18 is the actual funder's payment or an equivalent annuity representation, not a second charge on the same funded capital.

Lemma N.19 (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 (309) 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 $\gamma_i \pi^0_{i,t+1}$. 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. □

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 $\widehat{J}_{i,t} = J_{i,t}/c_t$. 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

$$\widehat{J}_{i,t} \geq \underline{J}_i := \delta_R(o_i - k_i - b_i) - \sum_{u=2}^{T_H+1} \delta_R^u \ell_{iu}. \tag{310}$$

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 N.20 (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*

$$S_V + \sum_i N_i s_{i,t} \leq S. \tag{311}$$

Choose a clean core $*$ and put $s_* = (S - S_V)/N_* > 0$. If

$$h_*'(s_*)[\underline{J}_* + G_*] - R_*'(s_*) > \max\left\{0, \max_{i \neq *}\left[h_i'(0)[\overline{J}_i + G_i]_+ - R_i'(0)\right]\right\}, \tag{312}$$

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.

The full proof is in the paper appendix, Section Q.8.1.

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.

N.4.4 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 $\widehat{R}_{a,t} = R_{a,t}/(w_r c_t)$ and $\widehat{b}_{a,t} = b_{a,t}/(w_r c_t)$. 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 $\widehat{\delta}_G = \delta_G g < 1$; at fixed expenditure $c_t = 1$ and $\widehat{\delta}_G = \delta_G < 1$. Normalized values obey

$$\begin{aligned} U_{a,t} &= \widehat{R}_{a,t} + \widehat{\delta}_G U_{a,t+1}, \\ V_{a,t} &= \max\{\widehat{b}_{a,t} + \widehat{\delta}_G V_{a,t+1}, -k_a + \widehat{\delta}_G[P_a U_{a,t+1} + (1-P_a)V_{a,t+1}]\}. \end{aligned} \tag{313}$$

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

$$\begin{aligned} \sigma_{a,t} &= u_{a,t} P_a, \\ M^D_{a,t+1} &= (1-\sigma_{a,t}) M^D_{a,t}, \qquad & M^C_{a,t+1} &= M^C_{a,t} + \sigma_{a,t} M^D_{a,t}. \end{aligned} \tag{314}$$

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 (303) 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 N.21 (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 prehistory. 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 (313) 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 $2\bar{r}\widehat{\delta}_G^{T-t+1}/(1-\widehat{\delta}_G)$, 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 (309), 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. □

N.4.5 Scarcity can strengthen protection: counterexample and repair

Proposition N.22 (Scarcity and a sufficient dominance condition)**.** *At a scalar final-good market of the form (304), 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

$$p_t \sum_j (\widehat{n}_{j,t+1} - \widehat{n}_{j,t}) Z_j(p_t) \geq e_{t+1} - e_t. \tag{315}$$

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

$$a(1-\rho) > (d_0 + \epsilon \max_a \omega_a^D) \max_{p \in [\underline{p}, \overline{p}]} p Z_D(p), \tag{316}$$

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 $F(p) = p \sum_j \widehat{n}_j Z_j(p) - e$. 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, (315) 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 (303) normalized clean mass is nondecreasing and normalized dirty mass satisfies $0 \leq \widehat{n}_{D,t} - \widehat{n}_{D,t+1} \leq (d_0 + \epsilon \max \omega_a^D) g^{-t}$, even if every ungoverned region intervenes. The normalized expenditure fall is $a(1-\rho)\rho^t$, at least that coefficient times g^{-t}. Condition (316) 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. □

Corollary N.23 (Dominance in the fully coupled infinite path)**.** *Use Lemma N.16. Let every reformable type install a strictly positive retail-clean component, so changes in b_s and B_D satisfy $|\Delta b_s| + |\Delta B_D| \le A\Delta\widehat{n}_C$ for a finite primitive ratio A. Keep e_c, e_d constant and let $e_{r,t} = e_0 + a\rho^t$, $g^{-1} \le \rho < 1$. There are finite uniform positive derivative bounds m_C, m_E, M_D, C obtained from the strict baseline market and Theorem N.13. If*

$$C\chi A < m_C, \qquad m_E a(1-\rho) > M_D(d_0 + \epsilon \max \omega_a^D) + C\chi(1 - g^{-1}), \tag{317}$$

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 $p^0_{\widehat{n}_C} < 0$, $p^0_{e_r} > 0$, and $p^0_{\widehat{n}_D} < 0$, 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 N.13, with χ small, preserves $p_{\widehat{n}_C} \le -m_C$, $p_{e_r} \ge m_E$, $|p_{\widehat{n}_D}| \le M_D$ and $|p_{b_s}| + |p_{B_D}| + |p_\lambda| \le C\chi$ 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

$$\Delta p \le -(m_C - C\chi A)\Delta\widehat{n}_C + M_D|\Delta\widehat{n}_D| - m_E|\Delta e_r| + C\chi|\Delta\lambda|.$$

Use $|\Delta\widehat{n}_D| \le (d_0 + \epsilon \max \omega_a^D)g^{-t}$, $|\Delta\lambda| = (1 - g^{-1})g^{-t}$ and $|\Delta e_r| = a(1-\rho)\rho^t$. Conditions (317) 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. □

The worked example is in Online Appendix O.11.

For physical extraction the relevant cell function is $q_D(p)$, not the transformed final-good quantity $Z_D(p)$. In the proved continuum-menu 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, (315), or the full coupled condition (317), 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.

N.4.6 When intervention becomes attractive

Theorem N.24 (Eventually strict intervention on the recleared path)**.** *In any connected equilibrium suppose each type has a funded legal-flow floor $\widehat{R}_{a,t} \geq \underline{R}_a > 0$ and normalized protection flows satisfy $0 \leq \widehat{b}_{a,t} \to 0$ uniformly over feasible population paths. If*

$$k_a(1 - \widehat{\delta}_G) < \widehat{\delta}_G P_a \underline{R}_a, \tag{318}$$

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.

The full proof is in the paper appendix, Section Q.10.1.

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 (318). All inequalities are open. The optional demand-drift condition strengthens this result to falling *absolute* dirty rents; growing legal opportunity is already sufficient 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 $-k - b_* + \widehat{\delta}_G P D_* > 0$. 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.

The worked example is in Online Appendix O.12.

N.4.7 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

$$T_{rh,t} = [\underline{y}_{rh} - y_{rh}(q_{r,t})]_+. \tag{319}$$

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 $\int T_{rh,t}\, dh \le w_r c_t f_{L,a}$. 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 $\pi^C_{a,t}/(w_r c_t) \ge \underline{\pi}_a > 0$, 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

$$0 < f_{S,a} + f_{L,a} < \tau_a \underline{\pi}_a, \qquad \underline{R}_a = \tau_a \underline{\pi}_a - f_{S,a} - f_{L,a} > 0. \tag{320}$$

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 N.25 (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,

$$\sum_a \bar{d}_a M^D_{a,0} \left(T + \frac{1}{P_a} \right) \le \mathcal{B}. \tag{321}$$

Successful legal production must itself remain in the specified safe implementation set. A certificate is not such a physical restriction.

The full proof is in the paper appendix, Section Q.11.1.

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.

N.4.8 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

$$c_r^g g_r = \alpha_m \mathcal{C}_r(q_r), \qquad H_r = (1-\alpha_m)\mathcal{C}_r(q_r), \qquad r_g q_r = \Pi_r + c_r^g g_r + H_r. \tag{322}$$

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 N.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 N.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 (274). 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.

For a continuum menu 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

$$\pi(p) = \max_{m,y\geq 0} \left\{ p[V(m) + H(y)] - \frac{k}{2}(m+y)^2 \right\}, \qquad q = m + y. \tag{323}$$

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.

N.5 Proof of Proposition N.26: Exact geographic aggregation

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 $\beta_r \, d\mu_g^{\mathrm{reg}}/B_g$ 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 N.1, Propositions N.6 and N.7. □

N.6 Proof of Theorem N.27: Joint common-price equilibrium

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 N.2. □

N.7 Proof of Proposition N.28: An exact active nonlinear class

Proof. Lemma N.10 establishes the unique global continuum-menu solution, positive retained deliveries, an active rationing premium, and positive resale. Equations (293) and (294) establish smooth physical output and the strict derivative signs for production, final output, and actual profit. The strict global comparison gaps and nonsingular Hessian persist under the finite-dimensional perturbations specified in Appendix N.3, giving the stated nonempty open parameter class.

Lemma N.11 proves the separate finite VCG construction, including attained global production, positive expected retained mineral, constant physical output on its strict branch, and increasing actual profit. Its proof also establishes that the efficient truthful allocation has no active terminal reallocation.

Theorem N.12 establishes the unique clearing price under the endpoint signs, continuity in type masses and demand, attained measurable household reconstruction, and the exact regional production and profit formulas. The positive core keeps the clearing derivative strictly positive when an individual type's mass vanishes. These results establish every claim of Proposition N.28. □

N.8 Retained baseline market definitions

The following source-model account states the definitions used by the retained baseline proofs; the revised paper states the scope of its newer results.

N.9 Regional production and world supply

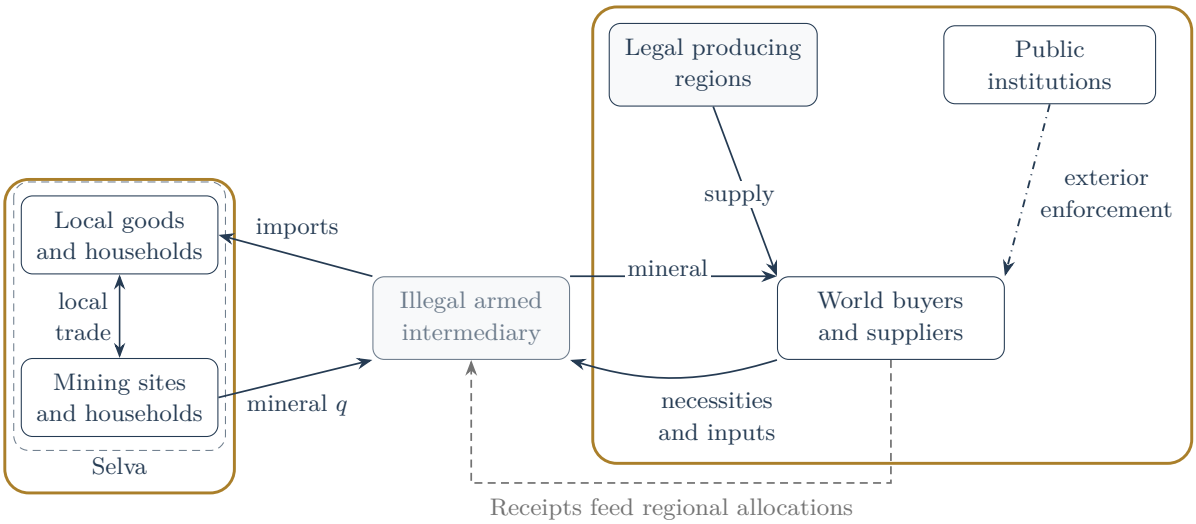

Regional supply enters world markets, while receipts feed back through regional allocations.

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 $(\mathcal{R}_g, \mu_g^{\mathrm{reg}})$ of regions, with mass $m_g = \mu_g^{\mathrm{reg}}(\mathcal{R}_g)$. 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

$$\mathcal{C}_r(q; A_r) = \frac{\eta}{1+\eta}\kappa_r A_r^{-1/\eta} q^{1+1/\eta}, \qquad \beta_r = A_r \kappa_r^{-\eta}, \qquad B_g = \int_{\mathcal{R}_g} \beta_r \, d\mu_g^{\mathrm{reg}} = m_g \mathbb{E}_g[\beta_r]. \tag{324}$$

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 N.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 $\max_q\{W_r(q;\Omega) - \mathcal{C}_r(q;A_r)\}$.

Proposition N.26 (Exact geographic aggregation)**.** *Under these conditions, output, operator profit and the least private cost of implementing group output are*

$$Q_g = B_g[r_g(\Omega)]_+^{\eta}, \qquad \Pi_g = \frac{B_g[r_g(\Omega)]_+^{1+\eta}}{1+\eta}, \tag{325}$$
$$\overline{\mathcal{C}}_g(Q) = \frac{\eta}{1+\eta} B_g^{-1/\eta} Q^{1+1/\eta}, \qquad q_r(Q) = \frac{\beta_r}{B_g} Q.$$

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. See Appendix N.5.

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.

N.9.1 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

$$Q_j = p_j^{\eta} \sum_{g \in j} B_g \omega_g^{\eta}, \qquad \Pi_j = \frac{p_j^{1+\eta}}{1+\eta} \sum_{g \in j} B_g \omega_g^{1+\eta}. \tag{326}$$

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 (326) characterize the multiplicative-receipt branch.

N.9.2 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

$$c_r^g g_r = \alpha_m \mathcal{C}_r(q_r), \qquad H_r = (1-\alpha_m)\mathcal{C}_r(q_r), \qquad r_g q_r = \Pi_r + c_r^g g_r + H_r. \tag{327}$$

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 N.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 N.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 (274). 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.

N.10 Certification and world-market equilibrium

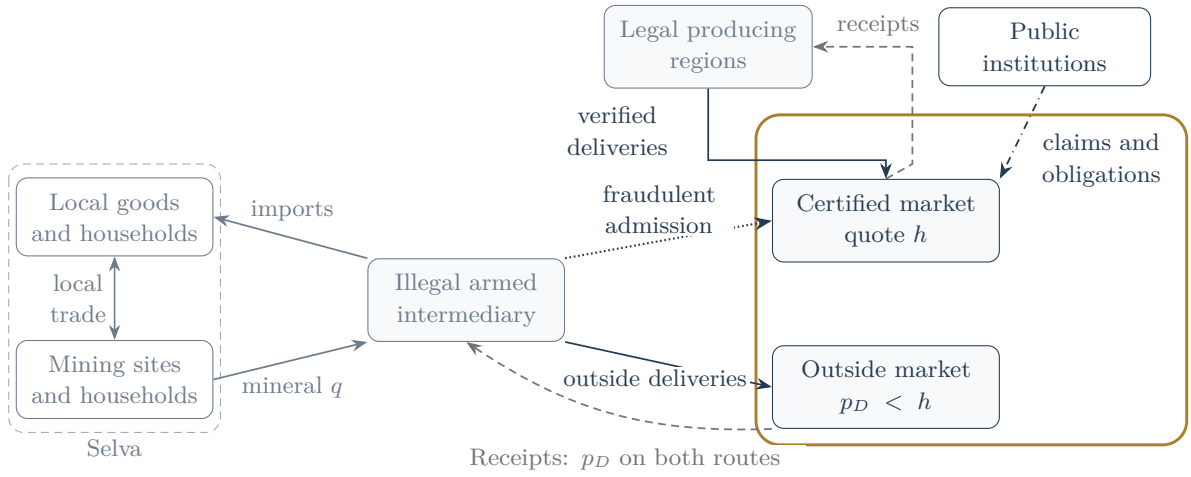

Certification separates markets; prices transmit its effects to regional production.

N.10.1 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 L. 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 M.

N.10.2 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

$$a_k = \frac{r_k l_{Lk} - l_{Ek}}{\kappa_k}, \qquad T_k(x) = 1 - x - r_k b_k, \qquad D_k = c_k + b_k(1 - r_k)a_k. \tag{328}$$

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

$$\pi_k(x) = \frac{T_k(x)}{D_k}, \quad e_k = a_k \pi_k, \quad p_k = r_k + (1 - r_k)e_k, \quad R_k = \pi_k r_k l_{Lk} - \frac{(r_k l_{Lk} - l_{Ek})^2 \pi_k^2}{2\kappa_k}. \tag{329}$$

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.

N.10.3 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

$$X_k = \frac{B_{Ck} h^\eta}{1 - \pi_k}, \qquad D_k^{\mathrm{div}} = \pi_k X_k, \qquad Y_k = Q_{Dk} - D_k^{\mathrm{div}}. \tag{330}$$

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:

$$\sum_k (Q_{Ck} + Q_{Dk}) = X + Y, \qquad Y = \sum_k Y_k. \tag{331}$$

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.

N.10.4 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

$$F_{\mathrm{cw}}(x) = \sum_k \frac{B_{Ck}}{1 - \pi_k(x)}, \qquad H_{\mathrm{cw}}(x) = F_{\mathrm{cw}}(x) - \sum_k B_{Ck}, \tag{332}$$

$$T(x) = \sum_k \frac{B_{Ck}[1 + \phi + R_k(x)]}{1 - \pi_k(x)}, \qquad J(x) = \frac{x[B_D x^\eta - H_{\mathrm{cw}}(x)]}{T(x)}. \tag{333}$$

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

Theorem N.27 (Joint common-price equilibrium)**.** *Suppose the individual incentive, resource and liquidity conditions above hold and $J(\underline{x}) < E_D/E_C < J(\overline{x})$. There is a unique equilibrium in this contract branch. It is characterized by*

$$J(x^*) = E_D/E_C, \qquad h^* = \left[\frac{E_C}{T(x^*)}\right]^{1/(1+\eta)}, \qquad p_D^* = x^* h^*. \tag{334}$$

All output, composition, fraud, effort and live claims follow (329)–(330). 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. See Appendix N.6.

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.

The worked example is in Online Appendix O.4.

N.10.5 Production rents and surviving contamination

The returns to clean and dirty producers are

$$\Pi_C = \frac{h^{1+\eta}}{1+\eta}\sum_k B_{Ck}, \qquad \Pi_D = \frac{(xh)^{1+\eta}}{1+\eta}\sum_k B_{Dk}. \tag{335}$$

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

$$\iota_t^{\text{live}} = \frac{\sum_{k,v}\mathbb{E}U_{kvt}}{\sum_{k,v}\mathbb{E}V_{kvt}}, \tag{336}$$

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 (349) 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.

N.11 Retail market power and the world price

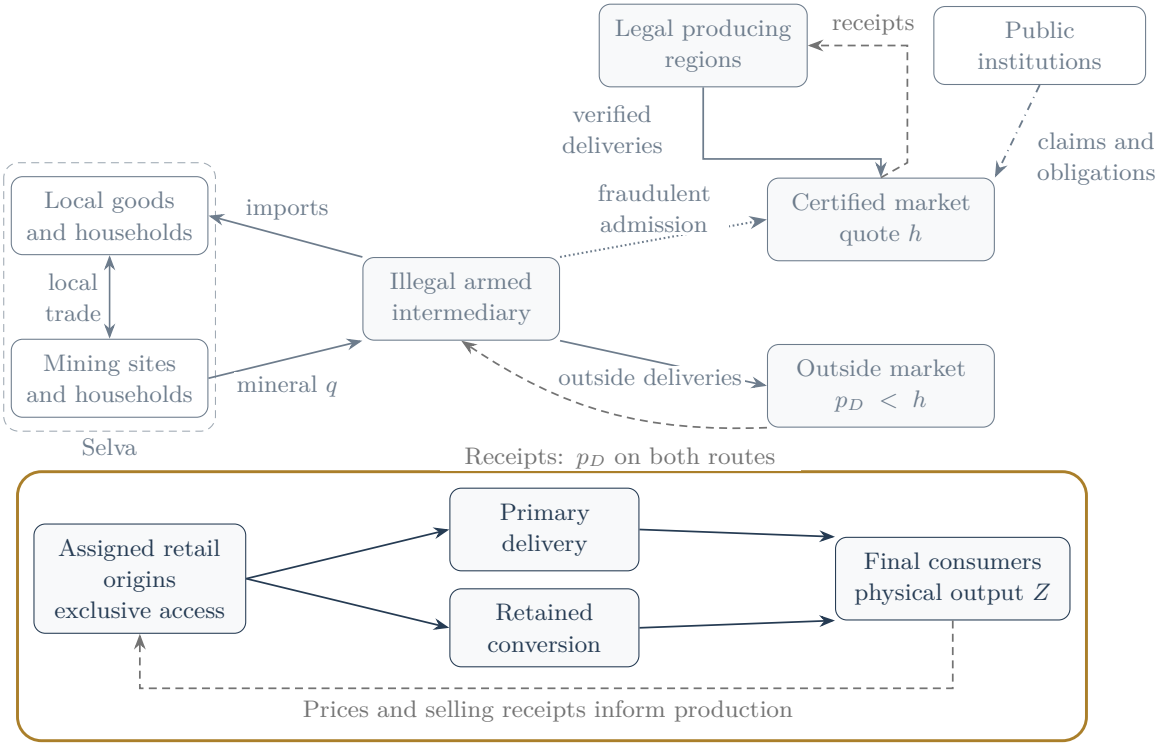

Exclusive processing routes connect local selling decisions to final demand and production receipts.

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.

The finite whole-buyer VCG mechanism has a nonlinear return to inventory and an efficient complete-game equilibrium with positive expected retention. To identify which production responses depend on indivisible buyer capacities, we also study a separate continuum extension. Its unit-demand buyers face a rationing menu, and its equilibrium has a positive rationing premium, resale and retained conversion. The extension changes the buyer population and selling institution; its active resale and smooth supply response are not properties of the finite VCG auction. Each regional production choice follows from its own optimized selling return.

N.11.1 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 continuum menu 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

$$\pi(p) = \max_{m,y\geq 0} \left\{ p[V(m) + H(y)] - \frac{k}{2}(m+y)^2 \right\}, \qquad q = m + y. \tag{337}$$

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

$$p \sum_j n_j Z_j(p) = E_R. \tag{338}$$

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 N.28 (An exact active nonlinear class)**.** *There is a nonempty open parameter class containing the explicit continuum-menu construction in Appendix N.3 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 (338) 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.

Proof. See Appendix N.7.

The continuum menu 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.

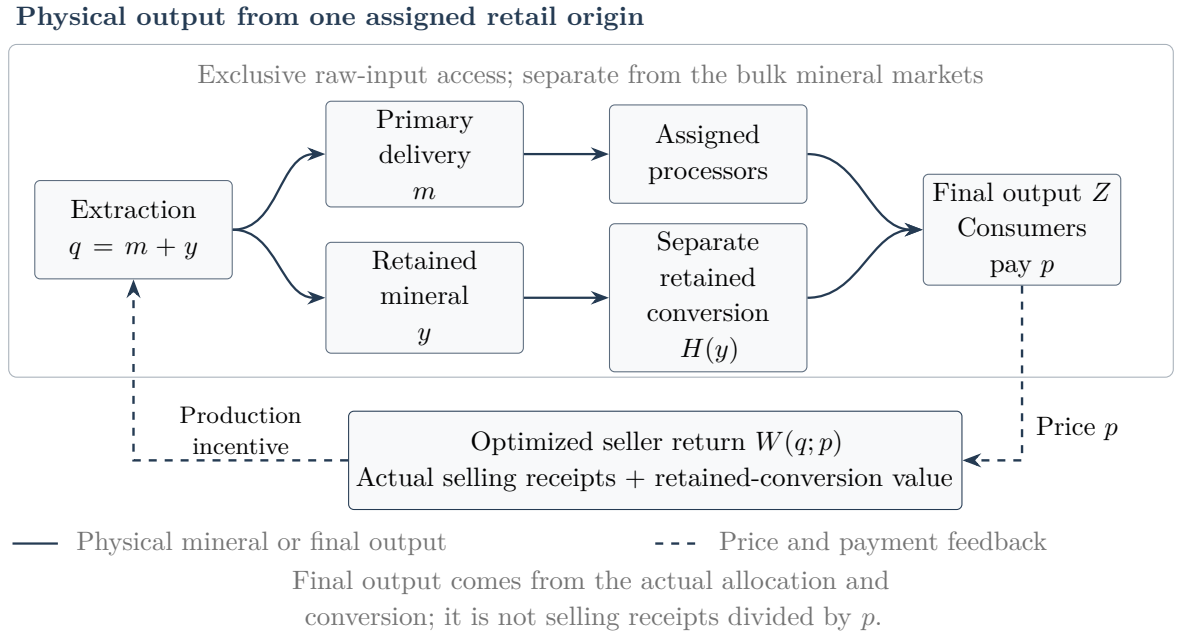

Selling branch	Physical response to higher p	Producer return
Continuum menu	$q'(p) > 0$ and $Z'(p) > 0$	Actual profit rises; the stated branch has active resale.
Finite VCG	Preferred inventory stays fixed on its strict branch	Expected profit rises; no active secondary trade is claimed.

Figure 3: 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 N.28.

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.

N.11.2 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

$$c_0 + E_C \log X + E_D \log Y + E_R \log Z_R + \varepsilon \log(Y + \gamma Z_R), \qquad \varepsilon, \gamma > 0. \tag{339}$$

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

$$\begin{aligned} hK(x)X &= E_C, \\ hxY &= E_D + \frac{\varepsilon Y}{Y + \gamma Z_R(p)}, \\ pZ_R(p) &= E_R + \frac{\varepsilon\gamma Z_R(p)}{Y + \gamma Z_R(p)}. \end{aligned} \tag{340}$$

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 N.29 (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 N.28, with positive retail core and endpoint clearing margins. For all sufficiently small $\varepsilon > 0$, (340) has a continuously differentiable equilibrium branch on a sufficiently small compact neighborhood of the baseline state. Its continuum-menu 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 proof is in Appendix N.3.5. 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 N.13 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 N.14 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.

N.11.3 The income available to innovators and governments

Consider balanced replication, with $n_{j,t} = L_t \widehat{n}_{j,t}$ and $E_{Rt} = L_t e_{Rt}$. Here $\widehat{n}_{j,t} = n_{j,t}/L_t$ is type j's cell mass per unit of scale and $\widehat{n}_t = (\widehat{n}_{j,t})_j$ is its vector. Dividing (338) by L_t leaves a price determined by $(\widehat{n}_t, e_{Rt})$, and type profit is $L_t \widehat{n}_{j,t} \pi_j(p_t)$. The same normalization applies to (340) 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 $\mathscr{W}_{\Sigma}(q - d \mid q, \Omega)$ 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.

N.12 Technical change and the incentive to govern

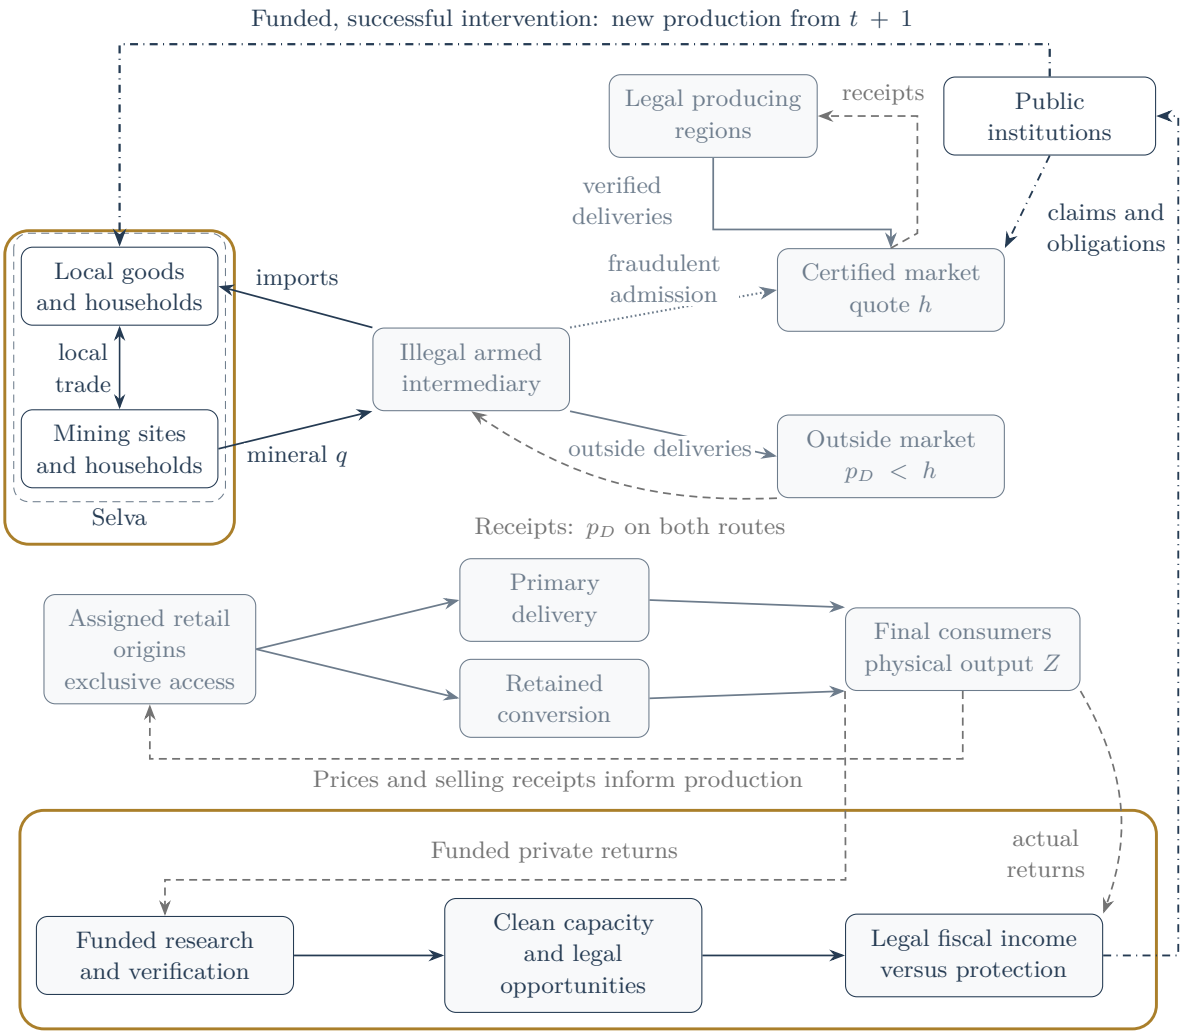

Funded investment and fiscal incentives can bring public authority into Selva, while funded livelihoods and safe production sites support protected-site conservation.

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.

N.12.1 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

$$W_t = a_0 Q_t + a_{\text{hist}} \sum_{s=1}^{T_H} Q_{t-s} + a_D D_t, \qquad D_{t+1} = D_t + n_t Q_t. \tag{341}$$

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

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

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 N.17 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 N.18 establishes its cash and resource feasibility. The evidence itself remains available after the financial obligation closes and after a region reforms.

N.12.2 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

$$S_V + \sum_i N_i s_{i,t} \le S. \tag{342}$$

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,

$$J_{i,t} = \text{discounted own success cash flow} - \text{discounted own failure cash flow}. \tag{343}$$

Equation (309) 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 N.19 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

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

Theorem N.20 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.

N.12.3 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 N.16 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.

N.12.4 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 $M_{a,t}^D = \int w_r 1_{\{r \text{ ungoverned}\}}\, d\mu_a^{\mathrm{reg}}$. If fraction $u_{a,t}$ intervenes and success has probability P_a, the exact population transition is

$$\begin{aligned} M_{a,t+1}^D &= (1 - u_{a,t}P_a)M_{a,t}^D, \\ M_{a,t+1}^C &= M_{a,t}^C + u_{a,t}P_a M_{a,t}^D. \end{aligned} \tag{344}$$

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 $\widehat{R}_{a,t}$ and $\widehat{b}_{a,t}$ 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 $\widehat{\delta}_G = \delta_G$ at fixed expenditure and $\widehat{\delta}_G = \delta_G g < 1$ in both growing-income classes. The values satisfy

$$\begin{aligned} U_{a,t} &= \widehat{R}_{a,t} + \widehat{\delta}_G U_{a,t+1}, \\ V_{a,t} &= \max\{\widehat{b}_{a,t} + \widehat{\delta}_G V_{a,t+1}, -k_a + \widehat{\delta}_G[P_a U_{a,t+1} + (1 - P_a)V_{a,t+1}]\}. \end{aligned} \tag{345}$$

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

Theorem N.21 establishes an infinite-horizon equilibrium of these government choices, research decisions and market prices. Its proof is in Appendix N.4.4. 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.

N.12.5 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

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

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 N.22 establishes this comparison, and Corollary N.23 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. Continuum-menu extraction 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.

N.12.6 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

$$k_a(1 - \widehat{\delta}_G) < \widehat{\delta}_G P_a \underline{R}_a \tag{346}$$

makes intervention optimal after a finite date for every region that remains ungoverned. Theorem N.24 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.

N.12.7 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 dis-

turbance 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 N.25 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.

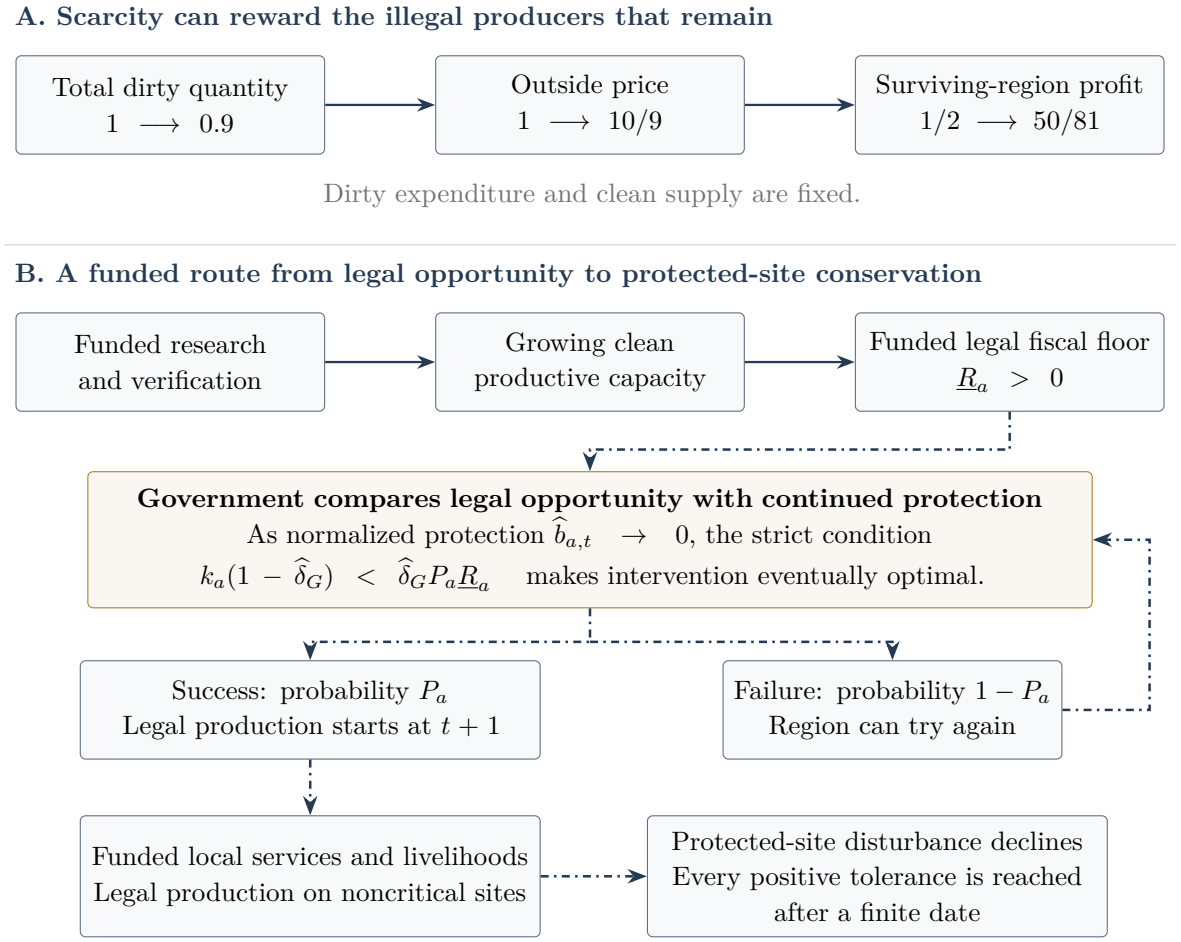

Figure 4: Scarcity and the conditional route to conservation. Panel A holds dirty expenditure and clean supply fixed: less aggregate illegal extraction can coexist with higher surviving-region rents. In panel B, navy dash-dot arrows show the conditional path and repeated attempts; the gold frame highlights 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.

O Worked examples and numerical material

O.1 Private and supported research scenarios

O.1.1 Private and supported research in two scenarios

Table 15 shows how private and supported research can give different industry paths within an already segmented market. The comparison holds government action off to isolate the research change. In gold, support raises clean effort to the selected allocation and lowers candidate output at period 100 by 10.7373% relative to the matched private case. In cobalt, private clean effort already reaches that allocation to numerical precision, and the supported rule pays no grants.

Table 15: Private and supported research: a direct model comparison.

	Candidate output, period 100		Paired change	Cumulative grants
Mineral	Private	Supported	(%)	dates 0–100
Gold	0.550439498	0.491337101	−10.7373	0.115314533
Cobalt	0.201390772	0.201390771	$\simeq 0$	0

The contrast makes support scenario dependent. The finite research rule evaluates incentives at the path's actual prices and keeps the reference demand-growth fit fixed across arms. It does not verify the infinite-horizon theorem's uniform sufficient condition. This direct comparison also differs from the joint contributions in Table 16, which average changes across orders of adding research support, government availability and the other factors.

O.2 Joint dynamic numerical comparisons

O.2.1 Joint changes in candidate extraction

The saved joint exercise shows how these opportunities enter the government's choice alongside research support, demand and discovery. Table 16 compares a common private-research, government-off reference with the complete four-factor package in segmented VCG markets. The output is candidate extraction at model period 100. Government action is available in the package, with the optimizing timing and success process, rather than imposed on every region.

Table 16: Contributions to the joint change in candidate extraction at model period 100.

Contribution (percentage points of baseline candidate extraction)	Gold	Cobalt
Outside-demand change	−14.7068	−11.4692
Stronger late discovery	−0.0094	−0.4163
Private versus supported research	−7.9564	$\simeq 0$
Government action available	−0.2590	−12.4928
Combined change	−22.9316	−24.3783

Notes: Segmented VCG, horizon 120. The common reference has private research, government action disabled, outside-spending cut 0 and late discovery 0.25. The complete package uses supported research, government action available, a 0.25 outside-spending cut approached over time and late discovery 0.50. Each contribution averages the factor's marginal change across all 24 orders in which the four factors can be added; all 16 required model cases are admitted for each mineral. Their sum is the combined change, up to rounding. Cobalt's research contribution is -2.26×10^{-7} percentage points. These are finite model comparisons; periods are not years and extraction is not forest loss. *Source:* Saved 10 October 2026 `arm_summary.csv` and `shapley.csv`; private-to-supported four-factor design, `qd100`.

Research support has a substantial contribution in the gold comparison, while government availability has a larger role in cobalt. These contributions average each factor's effect across the 24 orders in which the four factors can be added. Section 5's direct private-to-supported comparison holds the other factors fixed. Its gold VCG pair and this table share the same private/government-off baseline normalization, but measure different contrasts. Cobalt's near-zero research contribution is consistent with its selected research allocation requiring no grants.

The combined change does not imply that every instrument reinforces every other instrument. Gold's demand–research interaction is positive, while its research–government interaction is negative; the complete interactions and saved arms remain in the online appendix. These finite paths illustrate the joint mechanism without certifying the infinite-horizon theorem's uniform conditions. Their extraction outcome also leaves the physical conservation question to the funded implementation and safe-site conditions below.

O.3 Explicit finite whole-buyer cell

The explicit cell has two independent whole buyers with productivities uniform on $[0, 1]$, one unit package each. It uses the all-offered pivot auction and the complete terminal-consent protocol, with true retention productivity $\rho_V = 1/10$ and cost $k_V q^2/2$, $k_V = 1/5$. Whole-package capacities make the stock continuation nonlinear. On the strict price interval $[0.99, 1.01]$, its unique chosen inventory, final output and actual expected profit are

$$q_V = 1, \qquad Z_V = \frac{2 + \rho_V^3}{3}, \qquad \pi_V(p) = pR_V - \frac{k_V}{2}, \qquad R_V = \rho_V + \frac{(1 - \rho_V)^3}{3}. \quad (347)$$

Expected retention is $\rho_V^2 > 0$. The efficient truthful allocation attains the complete-game bound; terminal trade remains accessible, but no active secondary trade occurs on that

path. Within this strict interval higher p raises profit while chosen inventory and final output stay fixed. The separate continuum-menu alternative, with its active resale and positive output derivatives, is developed in the online appendix.

O.4 An equilibrium and its physical accounts

Example O.1 (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.

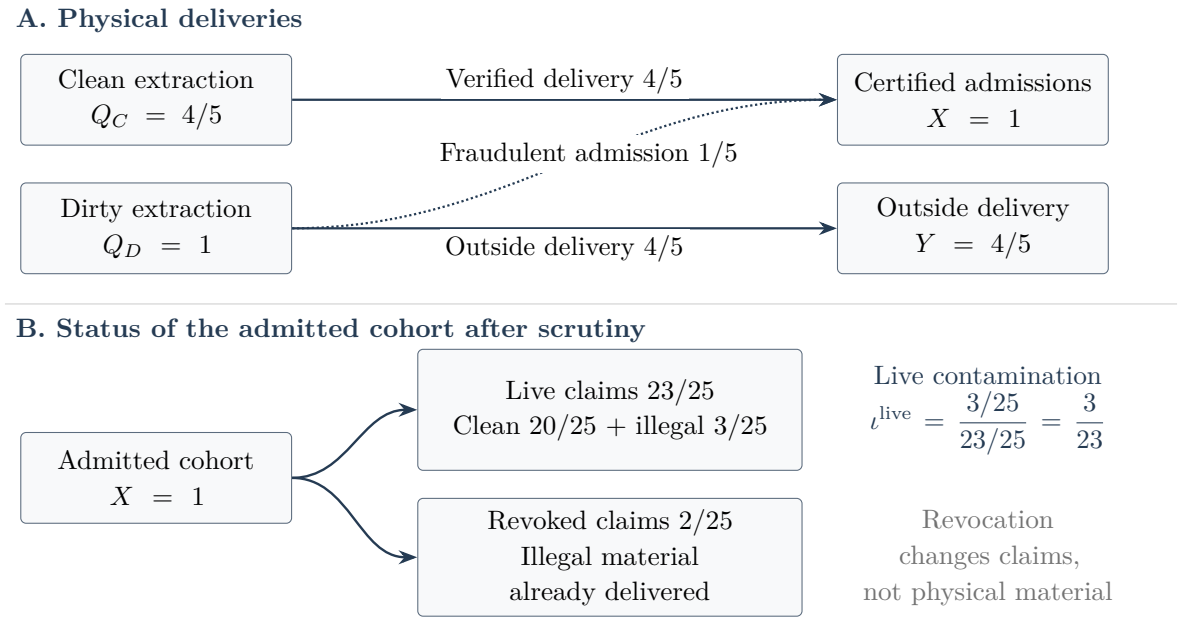

Figure 5: Deliveries and claim status in Example O.1. Panel A conserves mineral: total extraction and total delivery both equal 9/5; panel B follows the admitted unit through scrutiny, with $23/25 + 2/25 = 1$.

O.5 Aggregation and the distribution of market power

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 $Q = 0$.

O.6 Supply and profit use different moments

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.

O.7 A physical damage account beyond the supply moment

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

$$D = \alpha_m^2 t^s \int \frac{\delta_r \beta_r}{c_r^g}\, d\mu^{\mathrm{reg}} + t \int \zeta_r \beta_r \, d\mu^{\mathrm{reg}}. \tag{348}$$

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.

O.8 Entry costs and the power-envelope comparison

Equation (269) 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 $Q \le 1$ and $(1 + Q^2)/2$ for $Q \ge 1$. 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.

O.9 A minimum operating output and participation

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.

O.10 Checking the interior common-price equilibrium

At the parameters of Example O.1, $a = 1$, $D_k = 5/6$, $\pi(x) = 1 - 6x/5$ 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.

O.11 Scarcity can raise a remaining producer's rent

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.

O.12 Waiting before intervention becomes attractive

Waiting can be strictly optimal before intervention becomes attractive. For a primitive two-phase example take $\widehat{\delta}_G = 0.8$, $P = 0.5$, $k = 0.1$, $R_t \equiv 0.2$, 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.

O.13 Exact common-price equilibrium construction

O.13.1 An exact equilibrium and arbitrarily small positive leakage

The worked example is in Online Appendix O.10.

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

$$0 < \pi(x) = \frac{1 - x - rb}{c + b(1 - r)a} \leq \epsilon \tag{349}$$

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.

O.14 Observed evidence and saved numerical comparisons

O.14.1 Observed outcomes and construction

The observed evidence records mining activity and mine outcomes, together with trade destinations and export unit values. Colombia's unauthorized on-land alluvial-gold footprint rises from 58,909 hectares in 2018 to 79,631 in 2023, an increase of 35.2%. In the mapped degraded La Pampa area, the reported dredge counts rise from 148 in 2021 to 598 in August 2023, an increase of 304.1%. The latter evidence concerns renewed activity in already damaged land, alongside the source's account of reduced new deforestation after Operation Mercury. Area and dredge counts retain their own units; neither is converted into mineral tonnage. The accompanying direct-outcome table preserves the six Colombia observations and both La Pampa counts; see `data/observed/direct_illegal_mining_outcomes.csv`. Sources are UNODC/MME EVOA 2023, Figure 9, printed p. 49, and MAAP 193, 28 September 2023.

The published matched DRC comparison reports a 28 percentage-point reduction in recent FARDC mine presence (standard error 10.8 points, $p = .022$), with a 48% control mean. Annual cassiterite production has an estimated change of $-1,289.96$ kilograms per mine (standard error 11,599.002, $p = .913$), against a 26,377.98 kilogram control mean. The selected estimates cover 58 mines and 43 clusters surveyed in late 2019 in South Kivu and Maniema. They are external estimates from Table 12 of the November 2020 PRG/IPIS/SFR/Ulula report; the archive's 2021 name records dissemination. The matched comparison supplies evidence about interference while leaving the production effect imprecise. Its interpretation depends on selection into the matched sample. The complete published outcomes are in `data/observed/published_drc_policy_estimates.csv`; the source is the PRG/IPIS/SFR/Ulula November 2020 report, Table 12, printed p. 44.

Recorded destination shares divide each destination's import value by the sum of returned World-partner import values for the same product and year. Gold is unwrought gold (HS710812); cobalt is HS810520 and tin ores/concentrates HS260900. Mainland China's gold share changes from 25.78% in 2015 to 22.52% in 2024. The six Gulf Cooperation Council (GCC) destinations account for 8.88% in 2015 and 18.09% in 2023. Their 2024 total is unavailable because the UAE observation is missing. The supporting cobalt and tin China shares rise from 3.30% and 11.98% in 2005 to 70.42% and 76.72% in 2024. The route, matched-value/weight and mirror tables preserve checks on composition and reporting support. Recorded destinations are geographic aggregates, rather than a count of independent buyers setting the seller's payment. The frozen UN Comtrade construction and source-query trail are documented in `data/source_audits/D9_observed_evidence.md`; destination, route, value/weight and mirror tables are identified in the supporting-file index.

Colombia's monthly export unit-value exercise relates border USD per product-kilogram to a World Bank USD per fine-gold-kilogram benchmark. The 36-month 2019–2021 log association is 1.07079, with approximate interval using the reported heteroskedastic-

ity/autocorrelation (HAC) adjustment [0.83937, 1.30222]; its first-difference counterpart is 0.29370, with interval [−0.43418, 1.02158]. The twelve held-out 2022 observations give log root mean squared error (RMSE) 0.13164, compared with 0.06581 for the last-training-value benchmark. The primary prediction uses the realized contemporaneous gold benchmark. These descriptive relationships retain the product/metal distinction and their evaluation result. They do not identify the mine-level production response in the model.

O.14.2 Certification: rates, physical masses and saved sensitivity

The paper retains 15–20% as its working benchmark for illegal gold in current supply. The saved certification comparison has a different denominator: false-origin mass divided by live certified mass after withdrawals. The central gold VCG comparison is 3.469514 before and 1.465577 after stronger late discovery, per 10,000 live certified units. Late discovery changes from 0.25 to 0.50, while the saved equilibrium also adjusts fraudulent admissions and downstream effort. Table 17 reports the corresponding saved central comparisons for all five minerals and both institutions. Muir is the archived name for the continuum-menu cell described in Appendix N.3; VCG denotes the separate whole-buyer cell.

Table 17: False-origin units per 10,000 live certified units in the saved central comparisons

Mineral	Mechanism	Before / 10,000	After / 10,000
Gold	Muir	3.469514	1.465577
Gold	VCG	3.469514	1.465577
Cobalt	Muir	55.399268	24.564506
Cobalt	VCG	55.399268	24.564505
Tantalum	Muir	56.349866	24.306086
Tantalum	VCG	56.349866	24.306089
Tin	Muir	56.193583	24.237245
Tin	VCG	56.193583	24.237248
Tungsten	Muir	54.921520	23.677200
Tungsten	VCG	54.921520	23.677202

Note: Before and after refer to late discovery 0.25 and 0.50; the denominator is live certified mass after withdrawals.

The physical quantities in this paragraph are fractions of reference world mine production for gold. Candidate extraction is output from the model's declared candidate-origin strata; those strata and the assumed unknown-origin shares are not measurements of illegal production. In the central gold VCG comparison, false certified mass is 0.000132400665 before and 0.000055920798 after the change; live certified mass is 0.381611518 and 0.381561674. Candidate extraction is 0.491361177 and 0.491328469, while total extraction is 1 and

0.999993928. Dividing each false mass by its live certified mass yields the reported contamination rate. The proportion of earlier false mass remaining is a separate statistic; it must not be presented as a certified contamination rate.

The reference fixes bulk supply elasticity at one, unknown-origin candidate share at 0.5 and retail share at 0.25. Vulnerability is 10% at named eligible origins with positive backing and zero in the rest of the world. The saved main sensitivity varies unknown-origin candidate share over 0.1, 0.5 and 0.9, and retail share over 0.1, 0.25 and 0.5. There are 90 matched comparisons across the five minerals and two institutions, with five VCG withholding alternatives kept separately. These are declared assumption choices, without estimated probability weights. The full JSON, `data/certification/paper-reference-data.json`, retains each baseline and screening source identifier, all parameters, physical accounts and precise outcomes. The approximately 13.04% contamination in the complete worked example uses another parameter choice and remains illustrative.

O.14.3 Dynamic scenarios, availability and finance

The saved design contains 480 factorial states: five minerals, two mechanisms, two market regimes, three research regimes, two government-availability choices, two outside-demand paths and two late-discovery policies. One-at-a-time comparisons add 260 unique paths after repeated central or clipped cases are identified. Horizon refinements add 40, giving 780 unique paths. Model periods have no calendar-year interpretation. The archived declaration precedes the saved results; the reference fits are held fixed across policy arms.

The outside-demand switch changes spending along the declared path, and the discovery switch raises late discovery from 0.25 to 0.50. In the principal four-factor comparison the research switch moves from private to supported research; frozen research is a separate ablation. Government availability permits an optimizing action rather than imposing success. The pooled comparison retains historical liabilities while removing current certification obligations and pooling current bulk expenditure. It is a separate institutional setup, including supported-research cases. Research funding as such is therefore not described as unavailable in every pooled market.

Table 18 counts every declared path and preserves all unavailable outcomes. Numerically admitted paths have the saved status `admitted_finite_equilibrium`. A solver failure or runtime limit does not by itself establish equilibrium nonexistence. The archive includes its narrow independent nonexistence diagnoses for two exact frozen, government-off VCG stages; those diagnoses are not generalized to other unavailable paths.

Table 18: Declared, numerically admitted and unavailable paths by mineral and mechanism

Mineral	Mechanism	Declared	Admitted	Unavailable
Cobalt	Muir	78	75	3
Cobalt	VCG	78	60	18

Mineral	Mechanism	Declared	Admitted	Unavailable
Gold	Muir	78	77	1
Gold	VCG	78	54	24
Tantalum	Muir	78	55	23
Tantalum	VCG	78	40	38
Tin	Muir	78	56	22
Tin	VCG	78	43	35
Tungsten	Muir	78	59	19
Tungsten	VCG	78	36	42

Table 19 reports contributions to candidate extraction at model period 100, with saved horizon 120, in the segmented market, divided by the same private-research, government-off, demand-and-discovery-off reference for each mineral and mechanism. They average the contribution of a factor over all policy orders. The full four-factor attribution is shown only when all sixteen required paths are admitted. Paired effects in the accompanying arm table, `data/dynamic/arm_summary.csv`, hold the other scenario choices fixed and answer a different question. The research comparison belongs with Section 5's mechanism; the joint comparison belongs with Section 6.

Table 19: Contributions to candidate extraction at model period 100, in percentage points of the common reference

Mineral	Mechanism	D	S	R	G	Combined
Gold	Muir	-14.549164	-0.009128	-5.631977	-0.266197	-20.456466
Gold	VCG	-14.706783	-0.009433	-7.956431	-0.258963	-22.931610
Cobalt	Muir	-11.472458	-0.414759	-0.000000	-12.497413	-24.384630
Cobalt	VCG	-11.469172	-0.416250	-0.000000	-12.492840	-24.378263
Tantalum	Muir	unavailable	unavailable	unavailable	unavailable	unavailable
Tantalum	VCG	unavailable	unavailable	unavailable	unavailable	unavailable
Tin	Muir	unavailable	unavailable	unavailable	unavailable	unavailable
Tin	VCG	unavailable	unavailable	unavailable	unavailable	unavailable
Tungsten	Muir	unavailable	unavailable	unavailable	unavailable	unavailable
Tungsten	VCG	unavailable	unavailable	unavailable	unavailable	unavailable

Note: Entries are percentage points of the own mineral–mechanism private-research, government-off, demand-and-discovery-off reference at model period 100. D changes outside spending along the declared path; S raises late discovery from 0.25 to 0.50; R changes private to supported research; G makes government action available. Contributions average marginal changes over all 24 policy orders. Each complete decomposition requires

all sixteen paths. The combined column gives the change with all four factors on. The declared spending path and scenario switches are recorded in data/dynamic/design.json and data/dynamic/README.txt.

For gold VCG, the demand–research pair interaction is +5.630222 percentage points and the research–government interaction is −0.074372 points of the reference. The interaction table, data/dynamic/interactions.csv, retains every nonempty subset of factors; interactions are not automatically reinforcing. Cobalt's supported-minus-private research contribution is approximately zero in this decomposition. These saved comparisons are consistent with the theoretical allowance for zero grants, rather than a universal need for support.

The scenario files retain cumulative research spending, grants, installations, tax receipts, intervention costs, central funding bridges, custody principal, service reserves and bond principal. Appendix N.4 supplies the analytical workload and budget arguments, including open cohorts, permanent records and separate procurement finance. The complete arm, country-intervention, horizon, segmentation and one-at-a-time tables accompany this appendix. Cumulative quantities sum model dates 0 through 100 inclusive. Financial transfers retain their cash-account treatment and are not added again as real resource use.

O.14.4 Reproduction and statement correspondence

Supporting-file index The links below identify the supplied editable tables and source records. Paths are relative to the package root; links in the packaged PDF resolve from its `output/pdf/` folder. These files report summaries. Full date-by-date path files and engines remain in the original reproduction project identified by the replay instructions.

Activity observations data/observed/direct_illegal_mining_outcomes.csv.

Published DRC outcomes data/observed/published_drc_policy_estimates.csv.

Destination series and groups data/observed/global_destination_series.csv; data/observed/gold_requested_markets_groups.csv.

Trade checks data/observed/route_decomposition.csv; data/observed/matched_value_weight.csv; data/observed/mirror_comparison.csv.

Export unit-value fit and evaluation data/observed/gold_coefficients.csv; data/observed/gold_evaluation_metrics.csv; data/observed/gold_monthly_sample.csv.

Certification inputs and sensitivity data/certification/paper-reference-data.json; data/certification/leakage_correction.md.

Dynamic arms and attribution data/dynamic/arm_summary.csv; data/dynamic/shapley.csv; data/dynamic/interactions.csv; data/dynamic/support.csv.

Finance, horizon and sensitivity summaries data/dynamic/country_interventions.csv; data/dynamic/horizon_comparison.csv; data/dynamic/segmentation_comparison.csv; data/dynamic/oat_summary.csv.

Methods, declarations and replay data/source_audits/D9_observed_evidence.md; data/source_audits/D10_certification_calibration.md; data/source_audits/D11_dynamic_quantification.md; data/source_audits/selected_evidence.json; data/dynamic/design.json; data/dynamic/reference_fits.json; data/dynamic/README.txt; data/dynamic/OUTPUT_SCHEMA.txt; data/dynamic/REPLAY.txt; data/source_manifest.json.

Direct research pair locators. The private-to-supported research pairs in data/dynamic/arm_summary.csv are gold Muir `395713329251ee15` to `eaeb7cba2a1e1dbc`; gold VCG `4151094d8469c139` to `b87d8b1d4b025f93`; cobalt Muir `921ef80ae89dcc4e` to `5b3e9d012543292c`; and cobalt VCG `658f8801a57cecdb` to `cc0f7fde7dd83279`. These pairs use the segmented market, government action off, late discovery 0.25, outside-spending cut zero, horizon 120 and outcome period 100. They support the direct research comparison, separately from the four-factor contributions. The bundled arm table gives outcomes and finance summaries; complete date-by-date paths remain in the external archive identified by data/dynamic/REPLAY.txt.

The accompanying data manifest pins every copied source by SHA-256. All copied tables and the certification JSON are byte-identical to their sources. The appendix is a report of completed saved exercises, with no parameter fitting, solver rerun, new observation or stock-grace simulation. The existing full replay instructions are included as a source record; a scientific replay works in a separate copy and preserves the declarations, original failures and recovery provenance. Full path archives, engines and source snapshots remain in the original reproduction project and are identified by those instructions; this draft's table bundle alone is not a portable full engine release.

The formal proof package remains separately versioned. Its pilot checks parts of two result interfaces and certifies zero complete manuscript results. The proof register therefore retains all 74 source proof blocks. Numerical admission and model-period comparisons do not certify the infinite-horizon hypotheses of eventual intervention or safe-site conservation. The limited launch-stock explanation adds no equations, scenarios or formalization obligations. Existing checks of mine admission, inherited recycling history, accounting and statement correspondence remain necessary.

P Core economic proofs

Supporting technical results and all worked examples are in the Online Appendix. The following proofs use the assumptions stated with their results; references identify the complete supporting arguments.

P.1 Subsistence income and effort

P.1.1 The subsistence labor response

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. □

P.1.2 Consolidated resources

Proof. Sum the binding household budgets. The uncancelled 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 (10). Substituting $p_v^z = c_v^z + \mu_z$ and $p_v^g = c_v^g + \mu_g$ into (9) proves (11). The physical transport loss $\sum_{vw}(1 - 1/\delta_{vw})t_{vw}$ remains real even though its counterpart in local resident payments cancels. □

P.1.3 The subsistence financing floor

Fix an active network satisfying the resource-account assumptions of Proposition 1.4. Let its imported-subsistence deficit be $d_A = \sum_{v\in A} \alpha c_v^z c_v - \sum_{v\in A} a_v > 0$. Suppose other resource costs, retained cash and coercion costs are nonnegative, and downstream value obeys $W(q) \le \bar{v}q$, with $\bar{v} > 0$. Every implementation that earns nonnegative profit satisfies $q \ge d_A/\bar{v}$. This financing bound is separate from the interior labor-response result above.

Proof. The consolidated account and $C_v > c_v$ imply $C(z) \ge \sum_{v\in A} \alpha c_v^z c_v - \sum_{v\in A} a_v = d_A$. Nonnegative profit gives $W(q) \ge C(z)$, whereas the value ceiling gives $W(q) \le \bar{v}q$. Thus $d_A \le \bar{v}q$; division by the positive $\bar{v}$ proves the bound. This argument does not assert existence of an implementation or attainment of the bound. Those questions are addressed by the separate regional implementation results online. □

P.2 Different prices, the same output and profit

P.2.1 Attained effective terms, nonunique gross quotes

The following result uses the implementation and contract assumptions of Online Appendix A.1.

Proposition P.1 (Attained effective terms, nonunique gross quotes)**.** *With bounded coercion, finite outside utility and continuous nonnegative transport and coercion costs, every nonempty cost sublevel in* (9) *is compact in physical allocations and effective terms*

(w, B). *Every nonempty output fiber attains its cost minimum. Its nominal implementation set contains*

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

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

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

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. □

P.2.2 Regional output as an implemented choice

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 objective in (13). 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. □

P.3 Buffers against selective regulation

P.3.1 The opportunity floor

Proof. The unselected seller orders available when winner i is removed include the consecutive returns $\rho_1(y), \ldots, \rho_{k_i^*}(y)$. 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. □

P.3.2 The remaining-market floor

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 nth order of that sub-book. Including retention orders can only raise the unaffected-book floor. Combine this observation with (24) and divide the package price by $\bar{q}$ to compare it with θ_B. □

P.3.3 Global output resilience

Proof. Apply (32) separately to the two choices and subtract. Their implementation costs have no direct policy term. The resulting identity is (35); 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. □

P.3.4 Protection by unaffected payment-setting orders

Proof. The protected baseline loses no receipts. Every competitor's receipts weakly decrease by (31), 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. □

P.3.5 Falling rents with unchanged extraction

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. □

P.4 Partial transmission

P.4.1 Local production transmission

Proof. Differentiate all the regional first-order conditions: $\mathcal{H}z_\tau + W_{q\tau}a = 0$. Multiply by a^T to obtain (37). Negative definiteness gives $a^\mathsf{T}\mathcal{H}^{-1}a < 0$. 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 (38) and the comparison. □

P.4.2 An analytical partial-transmission condition

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 (39). Concavity makes $\rho' \leq 0$. The remaining conditions allow application of Theorem 2.14. □

P.5 Ceilings on receipt, profit and production losses

P.5.1 The exact floor and maximum loss of receipts

The following result uses the implementation and contract assumptions of Online Appendix A.2.

Proposition P.2 (The exact floor and maximum loss of receipts)**.** *For every feasible fixed stock and offered capacity,*

$$\begin{aligned} H(q) \leq \underline{W}_n(q) \leq W_n(q,\tau) \leq W_n(q,0), \\ 0 \leq W_n(q,0) - W_n(q,\tau) \leq \min\{n\bar{q}\tau, M_n(q)\} \end{aligned} \tag{350}$$

Moreover,

$$\lim_{\tau\to\infty} W_n(q,\tau) = \underline{W}_n(q), \qquad \sup_{\tau\geq 0}\{W_n(q,0) - W_n(q,\tau)\} = M_n(q) \tag{351}$$

Thus $M_n(q)$ is the exact maximum loss of expected receipts, understood as a supremum when the type distribution has unbounded support.

For $n \geq 1$, let $B_{\mathcal{R}}(\omega) = \max_{i\in\mathcal{R}}(b_i^1)_+$, with value zero when $\mathcal{R}$ is empty. A sufficient statewise saturation charge is

$$T_n(q,\omega) = \frac{[B_{\mathcal{R}}(\omega) - e_{\mathcal{U},H}(q,n,\omega)]_+}{\bar{q}} \tag{352}$$

For every $\tau \geq T_n(q,\omega)$, seller value equals $R_n^{\mathcal{U}}(q,\omega)$. Set $T_0 = 0$. Consequently the expected floor is attained for all charges at least $\operatorname{ess\,sup}_\nu T_n(q,\omega)$ whenever this bound is finite. The simpler bound

$$\overline{T} = \frac{\operatorname{ess\,sup}_\nu B_{\mathcal{R}}(\omega)}{\bar{q}} < \infty \tag{353}$$

is sufficient simultaneously for every stock and every allowed capacity.

Proof. Fix q, n, ω. All regulated marginal orders are no larger than their respective first order. Once $\bar{q}\tau \geq B_{\mathcal{R}}(\omega)$, each has nonpositive value. Free disposal and nondecreasing H allow each such allocation to be replaced by retention without reducing welfare, both in the full problem and after excluding any buyer. Consequently those welfare values equal the corresponding values when regulated buyers are absent. The seller-value identity $R_n = \sum_i S_{n,-i} - (N-1)S_n$ then reduces to the unaffected-only identity: each absent regulated buyer contributes S_n, cancelling its contribution to the multiplier $N-1$. Thus seller value equals $R_n^{\mathcal{U}}$, including at zero-value ties.

The exact exposure identity makes $R_n(q,\omega,\tau)$ weakly decreasing in τ. It follows that it is bounded below by its eventual value $R_n^{\mathcal{U}}$. The opportunity floor applied to the surviving auction gives $R_n^{\mathcal{U}} \geq H(q)$. The original finite-stock integrability assumption bounds expected receipts: each truthful buyer pays no more than its own allocated value, and retained value is at most $H(q)$, so an integrable dominating variable is $H(q) + \sum_{i,k}(b_i^k)_+$. Dominated convergence therefore gives $W_n(q,\tau) \downarrow \underline{W}_n(q)$. Combining this bound with the integrated exposure bound $n\bar{q}\tau$ proves (350) and (351).

For the sharper cutoff, the unaffected-plus-retention book in every buyer-exclusion problem has at least n orders of value at least $e_{\mathcal{U},H}(q,n,\omega)$. The full book contains all these orders too. At a charge satisfying (352), every regulated order is weakly below that level. The highest n values in each relevant maximization can therefore all be supplied by unaffected buyers and retention. Ties leave the welfare value unchanged. The same seller-value identity gives exact equality to the surviving auction. Taking an essential supremum yields the fixed-stock expected cutoff; nonnegativity of the unaffected-plus-retention floor gives the common bound (353). When $n = 0$, receipts are $H(q)$ at every policy.

Finally, under the strict demand-depth premise preceding (33), every optional retention order is below at least n unaffected orders. The surviving auction therefore sells all n packages externally and retains precisely $q - n\bar{q}$. Each of its n critical-payment slots is at least $e_{\mathcal{U},H}$ by Proposition 2.9. Summing slots and taking expectations gives (33). □

P.5.2 Maximum profit loss and a quantity bound

The following result uses the implementation and contract assumptions of Online Appendix A.2.

Proposition P.3 (Maximum profit loss and a quantity bound)**.** *For every charge,*

$$0 \leq \Pi^*(0) - \Pi^*(\tau) \leq \Pi^*(0) - \underline{\Pi} \leq D_0 \tag{354}$$

If the uniform cutoff (353) *is finite, then* $\Pi^*(\tau) = \underline{\Pi}$ *for all* $\tau \geq \overline{T}$*, so* $\Pi^*(0) - \underline{\Pi}$ *is the exact maximum optimized-profit loss. At every policy, any selected optimal quantity* q_τ *satisfies*

$$\Delta_0(q_\tau) \leq B_0(\tau) \leq D_0 \tag{355}$$

If, on every feasible quantity including exit, baseline profit has the quadratic separation $\Delta_0(q) \geq \kappa(q - q_0)^2/2$ *for some* $\kappa > 0$*, then*

$$|q_\tau - q_0| \leq \sqrt{\frac{2B_0(\tau)}{\kappa}} \leq \sqrt{\frac{2D_0}{\kappa}} \tag{356}$$

Proof. For every quantity, including exit, $\underline{V}(q) \leq V(q,\tau) \leq V(q,0)$. Maximization gives $\underline{\Pi} \leq \Pi^*(\tau) \leq \Pi^*(0)$. Since $\underline{\Pi} \geq \underline{V}(q_0) = \Pi^*(0) - D_0$, the last inequality in (354) follows. Under the finite uniform cutoff, every allowed capacity at every quantity equals its surviving-market value. Their maxima consequently agree as well, giving $\Pi^*(\tau) = \underline{\Pi}$ at every later charge.

For any policy optimum, baseline value, policy value, optimality and the baseline-output loss certificate give the successive inequalities

$$V(q_\tau, 0) \geq V(q_\tau, \tau) \geq V(q_0, \tau) \geq V(q_0, 0) - B_0(\tau)$$

Rearrangement proves (355). Substituting the quadratic separation premise and taking square roots proves (356). At and beyond a finite uniform cutoff, the entire objective and hence its complete argmax set are constant. Every contraction occurring later is consequently already represented at the cutoff in (46). Nonnegative output and the quantity bound yield the displayed upper bound on $\mathcal{D}_q$. No ordering of the maximizing quantities across different policy levels was used or follows from these arguments. □

P.5.3 A resource bound on policy effects

The following result uses the implementation and contract assumptions of Online Appendix A.2.

Proposition P.4 (A resource bound on policy effects)**.** *Normalize exit profit to zero. Let* $\mathcal{Q}$ *be the fixed set of attainable positive outputs and let* $\mathcal{T}$ *be a policy interval containing zero. Suppose the policy profit maximum is finite and attained for every* $\tau \in \mathcal{T}$*. Assume uniform constants* $\bar{v} > 0$ *and* $d > 0$*, and a policy-independent attainable continuation floor* F*, such that*

$$F(q) \leq W(q,\tau) \leq \bar{v}q, \qquad \mathcal{C}(q) \geq d \quad (q \in \mathcal{Q},\ \tau \in \mathcal{T}) \tag{357}$$

Define the protected profit

$$\Pi_F = \max\left\{0, \sup_{q \in \mathcal{Q}}[F(q) - \mathcal{C}(q)]\right\} > 0 \tag{358}$$

Every policy optimum is active and satisfies

$$q_\tau \geq \underline{q}_F := \frac{d + \Pi_F}{\bar{v}} > 0 \tag{359}$$

Consequently, for any baseline optimum q_0 and any $\tau \in \mathcal{T}$,

$$\frac{(q_0 - q_\tau)_+}{q_0} \leq 1 - \frac{d + \Pi_F}{\bar{v} q_0} \tag{360}$$

The supremum defining Π_F need not be attained.

Proof. Taking suprema in $W(q,\tau) - \mathcal{C}(q) \geq F(q) - \mathcal{C}(q)$ and including exit gives optimized profit at least Π_F, without requiring the floor supremum to attain. The finite policy maximum also bounds Π_F. Its strict positivity excludes exit. Every optimum therefore satisfies $\Pi_F \leq W(q_\tau,\tau) - \mathcal{C}(q_\tau) \leq \bar{v} q_\tau - d$, proving (359). The same bound at zero policy gives $q_0 \geq \underline{q}_F$. Thus $(q_0 - q_\tau)_+ \leq q_0 - \underline{q}_F$; division by $q_0 > 0$ proves (360). □

P.5.4 Global policy ceiling

Proof. For each permitted capacity, Proposition P.2 gives $W_n(q,\tau) \geq \underline{W}_n(q)$. Maximizing both sides over the same capacity menu proves (42).

Take the limit as $\tau \to \infty$ in (32). Proposition P.2 identifies the limiting value as $\underline{W}_n(q)$. Since expected exposure is nonnegative, its improper integral exists and equals the finite receipt difference. It is positive exactly when $\bar{\chi}_n(q,t) > 0$ on a set of positive Lebesgue measure. Expected exposure records regulated orders in buyer-exclusion payment comparisons, including orders that lose in the actual allocation.

For any policy-optimal quantity q_τ, the remaining-market floor, optimality and weakly decreasing receipts give

$$V(q_\tau, 0) \geq V(q_\tau,\tau) = \Pi^*(\tau) \geq V(q_0,\tau) \geq V(q_0,0) - D_0$$

Exit is included with its unchanged payoff. These inequalities give the profit-loss and baseline-gap bounds. If $D_0 = 0$, baseline uniqueness forces $q_\tau = q_0$, and profit is unchanged. If every quantity outside the ε-neighborhood has gap greater than D_0, the gap bound excludes each such quantity.

Finally, if one baseline-optimal offered capacity n_0 at q_0 has zero cumulative exposure, then $\underline{W}_{n_0}(q_0) = W_{n_0}(q_0,0) = W(q_0,0)$. The optimized floor lies between this value and $W(q_0,0)$, so $D_0 = 0$. Thus positive exposure in every baseline-optimal selling choice is necessary for a production response. The protected marginal-return cases in Corollary A.21 show why it is not sufficient.

For the positive production floor, strict positivity of $\underline{\Pi}$ excludes exit. At any chosen active output,

$$\underline{\Pi} \leq W(q_\tau,\tau) - \mathcal{C}(q_\tau) \leq \bar{v} q_\tau - d$$

Rearranging gives (43). This is the resource bound of Proposition P.4 with $F = \underline{W}$. It is uniform over the full tariff path and also applies to the auction with all regulated buyers removed. The guaranteed floor need not itself be feasible or attained; restrictions on feasible quantities can raise the smallest possible chosen output. □

P.6 Physical and incentive architecture

P.6.1 Conservation and retrospective exclusion

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 (45). 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. □

P.6.2 Investigation and voluntary disclosure

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 (47). 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 (48). Its strict concavity establishes uniqueness; its first derivative gives (49). Implicit differentiation gives (69). 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. □

P.6.3 Funded capture with optional suppression

Proof. For a fixed $\mathcal{T}$, original recipient i gains $\lambda_i t_i - a_i(\mathcal{T})$ relative to honesty. Acceptance therefore requires $t_i \geq a_i(\mathcal{T})/\lambda_i$. Each purchased reporter analogously requires $t_j \geq \sigma_j(\mathcal{T})/\lambda_j$. The joint implementability assumption makes these bounds sufficient as well as necessary. Thus every accepted schedule implementing $\mathcal{T}$ costs at least $\mathcal{B}(\mathcal{T})$ before settlement. If that bill exceeds F_0, no such schedule is funded. Otherwise equality payments

are funded and leave organizer surplus $G_\Omega(a) - a_0(\mathcal{T}) - \mathcal{B}(\mathcal{T})$; 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. □

P.6.4 Positive-trade interior equilibrium

Proof. 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.

Strict convexity in (56) gives the buyer's unique best response. For an interior solution the uniform cost distribution gives $\bar{c}\pi = \Gamma - rD - D(1-r)a\pi$, yielding (58). The solution lies strictly inside the unit interval because its numerator is positive and less than $\bar{c}$; the last inequality in (57) ensures $e^* = a\pi^* < 1$. More generally, the fraud best response is the continuous, nonincreasing map $\pi \mapsto [\{\Gamma - DP(\min(a\pi, 1))\}/\bar{c}]_0^1$, 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 $p\sum_i b_i \le F$. The opening bargaining comparison establishes sequential reporting and the absence of mutually acceptable silence contracts. Posted wages, prepaid fees and segregated deposits satisfy the stated participation and budget contracts; buyer participation completes the positive-trade allocation.

For the matched benchmark, $\pi_0 = \widetilde{\Gamma}/\bar{c}$ and $e_0 = 0$. Equation (58) gives $\pi^* < \pi_0$ and $P^* > r$. The live-contamination function in (60) has derivatives $(1-P)/(1-\pi P)^2 > 0$ 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. □

P.7 The certification price shock

Proof. At the pooled price compare certified supply with the demand redirected to the certified segment. Strict monotonicity of the market-clearing equation gives the signs in the homogeneous benchmark. When ordinary labor and headquarters adjust, eliminate their first-order conditions and use the unique dirty-price crossing of the resulting market root. The full derivation and the distinction between gross and net receipts are in Online

Appendix E.2. No fixed expenditure assumption is carried from this date comparison into the growth path. □

P.8 Patent ownership, research and withdrawal

The production and ownership accounts in Online Appendix E.2 and E.6 construct the certified machine demand and the surviving vintage cash flow. Conditioning the survival sum on the next date gives the valuation recursion; its dominated tail gives the terminal condition. Dirty one-period advantage is a separate incremental production claim with subsequent diffusion and retained contractual settlement.

Proof of the research and path-dependence results. For given correctly valued rewards, the concave research problem has the stated marginal-return conditions, including inactivity. At the all-certified allocation these reduce to the threshold in Theorem I.8. Online Appendix E.7 proves necessity and sufficiency and handles ties. The private-green set must be invariant under unsubsidized equilibrium continuation; induction then permits permanent withdrawal. Bounds on certified and dirty rewards give an explicit sufficient region, and paid growth reaches it under the conditions in Online Appendix E.8. A one-date return comparison or an unfunded subsidy promise does not prove this result. □

P.9 Taxation, intervention and physical conservation

The fiscal objective uses the privately feasible company response and the same dated costs as the control value. Its tax elasticity therefore includes headquarters adjustment. The Bellman contraction and action-gap limit give Theorem J.1 on the stated feasible cash domain. Conditional failure products give individual success and population bounds; the physical disturbance bound then gives the distinct conclusions in Proposition J.2. Complete derivations, financing obstructions and the conditional conservation-game construction are in the current asymmetric proofs in the Online Appendix. The earlier government-only fixed point is not a construction of the expanded company, research and market economy.

Q Core economic proofs

Supporting technical results and all worked examples are in the Online Appendix. The following proofs use the assumptions stated with their results; references identify the complete supporting arguments.

Q.1 Subsistence income and effort

Q.1.1 The subsistence labor response

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. □

Q.1.2 Consolidated resources

Proof. Sum the binding household budgets. The uncancelled 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 (10). Substituting $p_v^z = c_v^z + \mu_z$ and $p_v^g = c_v^g + \mu_g$ into (9) proves (11). The physical transport loss $\sum_{vw}(1 - 1/\delta_{vw})t_{vw}$ remains real even though its counterpart in local resident payments cancels. □

Q.1.3 The subsistence financing floor

Fix an active network satisfying the resource-account assumptions of Proposition 1.4. Let its imported-subsistence deficit be $d_A = \sum_{v\in A} \alpha c_v^z c_v - \sum_{v\in A} a_v > 0$. Suppose other resource costs, retained cash and coercion costs are nonnegative, and downstream value obeys $W(q) \le \bar{v}q$, with $\bar{v} > 0$. Every implementation that earns nonnegative profit satisfies $q \ge d_A/\bar{v}$. This financing bound is separate from the interior labor-response result above.

Proof. The consolidated account and $C_v > c_v$ imply $C(z) \ge \sum_{v\in A} \alpha c_v^z c_v - \sum_{v\in A} a_v = d_A$. Nonnegative profit gives $W(q) \ge C(z)$, whereas the value ceiling gives $W(q) \le \bar{v}q$. Thus $d_A \le \bar{v}q$; division by the positive $\bar{v}$ proves the bound. This argument does not assert existence of an implementation or attainment of the bound. Those questions are addressed by the separate regional implementation results online. □

Q.2 Different prices, the same output and profit

Q.2.1 Attained effective terms, nonunique gross quotes

The following result uses the implementation and contract assumptions of Online Appendix A.1.

Proposition Q.1 (Attained effective terms, nonunique gross quotes)**.** *With bounded coercion, finite outside utility and continuous nonnegative transport and coercion costs, every nonempty cost sublevel in* (9) *is compact in physical allocations and effective terms*

(w, B). *Every nonempty output fiber attains its cost minimum. Its nominal implementation set contains*

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

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

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

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. □

Q.2.2 Regional output as an implemented choice

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 objective in (13). 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. □

Q.3 Buffers against selective regulation

Q.3.1 The opportunity floor

Proof. The unselected seller orders available when winner i is removed include the consecutive returns $\rho_1(y), \ldots, \rho_{k_i^*}(y)$. 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. □

Q.3.2 The remaining-market floor

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 nth order of that sub-book. Including retention orders can only raise the unaffected-book floor. Combine this observation with (24) and divide the package price by $\bar{q}$ to compare it with θ_B. □

Q.3.3 Global output resilience

Proof. Apply (32) separately to the two choices and subtract. Their implementation costs have no direct policy term. The resulting identity is (35); 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. □

Q.3.4 Protection by unaffected payment-setting orders

Proof. The protected baseline loses no receipts. Every competitor's receipts weakly decrease by (31), 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. □

Q.3.5 Falling rents with unchanged extraction

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. □

Q.4 Partial transmission

Q.4.1 Local production transmission

Proof. Differentiate all the regional first-order conditions: $\mathcal{H}z_\tau + W_{q\tau}a = 0$. Multiply by a^T to obtain (37). Negative definiteness gives $a^\mathsf{T}\mathcal{H}^{-1}a < 0$. 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 (38) and the comparison. □

Q.4.2 An analytical partial-transmission condition

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 (39). Concavity makes $\rho' \le 0$. The remaining conditions allow application of Theorem 2.14. □

Q.5 Ceilings on receipt, profit and production losses

Q.5.1 The exact floor and maximum loss of receipts

The following result uses the implementation and contract assumptions of Online Appendix A.2.

Proposition Q.2 (The exact floor and maximum loss of receipts)**.** *For every feasible fixed stock and offered capacity,*

$$\begin{aligned} H(q) \le \underline{W}_n(q) \le W_n(q,\tau) \le W_n(q,0), \\ 0 \le W_n(q,0) - W_n(q,\tau) \le \min\{n\bar{q}\tau, M_n(q)\} \end{aligned} \tag{361}$$

Moreover,

$$\lim_{\tau\to\infty} W_n(q,\tau) = \underline{W}_n(q), \qquad \sup_{\tau\ge 0}\{W_n(q,0) - W_n(q,\tau)\} = M_n(q) \tag{362}$$

Thus $M_n(q)$ is the exact maximum loss of expected receipts, understood as a supremum when the type distribution has unbounded support.

For $n \ge 1$, let $B_{\mathcal{R}}(\omega) = \max_{i\in\mathcal{R}}(b_i^1)_+$, with value zero when $\mathcal{R}$ is empty. A sufficient statewise saturation charge is

$$T_n(q,\omega) = \frac{[B_{\mathcal{R}}(\omega) - e_{\mathcal{U},H}(q,n,\omega)]_+}{\bar{q}} \tag{363}$$

For every $\tau \geq T_n(q,\omega)$, seller value equals $R_n^{\mathcal{U}}(q,\omega)$. Set $T_0 = 0$. Consequently the expected floor is attained for all charges at least $\operatorname{ess\,sup}_\nu T_n(q,\omega)$ whenever this bound is finite. The simpler bound

$$\overline{T} = \frac{\operatorname{ess\,sup}_\nu B_{\mathcal{R}}(\omega)}{\bar{q}} < \infty \tag{364}$$

is sufficient simultaneously for every stock and every allowed capacity.

Proof. Fix q, n, ω. All regulated marginal orders are no larger than their respective first order. Once $\bar{q}\tau \geq B_{\mathcal{R}}(\omega)$, each has nonpositive value. Free disposal and nondecreasing H allow each such allocation to be replaced by retention without reducing welfare, both in the full problem and after excluding any buyer. Consequently those welfare values equal the corresponding values when regulated buyers are absent. The seller-value identity $R_n = \sum_i S_{n,-i} - (N-1)S_n$ then reduces to the unaffected-only identity: each absent regulated buyer contributes S_n, cancelling its contribution to the multiplier $N-1$. Thus seller value equals $R_n^{\mathcal{U}}$, including at zero-value ties.

The exact exposure identity makes $R_n(q,\omega,\tau)$ weakly decreasing in τ. It follows that it is bounded below by its eventual value $R_n^{\mathcal{U}}$. The opportunity floor applied to the surviving auction gives $R_n^{\mathcal{U}} \geq H(q)$. The original finite-stock integrability assumption bounds expected receipts: each truthful buyer pays no more than its own allocated value, and retained value is at most $H(q)$, so an integrable dominating variable is $H(q) + \sum_{i,k}(b_i^k)_+$. Dominated convergence therefore gives $W_n(q,\tau) \downarrow \underline{W}_n(q)$. Combining this bound with the integrated exposure bound $n\bar{q}\tau$ proves (361) and (362).

For the sharper cutoff, the unaffected-plus-retention book in every buyer-exclusion problem has at least n orders of value at least $e_{\mathcal{U},H}(q,n,\omega)$. The full book contains all these orders too. At a charge satisfying (363), every regulated order is weakly below that level. The highest n values in each relevant maximization can therefore all be supplied by unaffected buyers and retention. Ties leave the welfare value unchanged. The same seller-value identity gives exact equality to the surviving auction. Taking an essential supremum yields the fixed-stock expected cutoff; nonnegativity of the unaffected-plus-retention floor gives the common bound (364). When $n = 0$, receipts are $H(q)$ at every policy.

Finally, under the strict demand-depth premise preceding (33), every optional retention order is below at least n unaffected orders. The surviving auction therefore sells all n packages externally and retains precisely $q - n\bar{q}$. Each of its n critical-payment slots is at least $e_{\mathcal{U},H}$ by Proposition 2.9. Summing slots and taking expectations gives (33). □

Q.5.2 Maximum profit loss and a quantity bound

The following result uses the implementation and contract assumptions of Online Appendix A.2.

Proposition Q.3 (Maximum profit loss and a quantity bound)**.** *For every charge,*

$$0 \leq \Pi^*(0) - \Pi^*(\tau) \leq \Pi^*(0) - \underline{\Pi} \leq D_0 \tag{365}$$

If the uniform cutoff (364) *is finite, then* $\Pi^*(\tau) = \underline{\Pi}$ *for all* $\tau \geq \overline{T}$*, so* $\Pi^*(0) - \underline{\Pi}$ *is the exact maximum optimized-profit loss. At every policy, any selected optimal quantity* q_τ *satisfies*

$$\Delta_0(q_\tau) \leq B_0(\tau) \leq D_0 \tag{366}$$

If, on every feasible quantity including exit, baseline profit has the quadratic separation $\Delta_0(q) \geq \kappa(q - q_0)^2/2$ *for some* $\kappa > 0$*, then*

$$|q_\tau - q_0| \leq \sqrt{\frac{2B_0(\tau)}{\kappa}} \leq \sqrt{\frac{2D_0}{\kappa}} \tag{367}$$

Proof. For every quantity, including exit, $\underline{V}(q) \leq V(q,\tau) \leq V(q,0)$. Maximization gives $\underline{\Pi} \leq \Pi^*(\tau) \leq \Pi^*(0)$. Since $\underline{\Pi} \geq \underline{V}(q_0) = \Pi^*(0) - D_0$, the last inequality in (365) follows. Under the finite uniform cutoff, every allowed capacity at every quantity equals its surviving-market value. Their maxima consequently agree as well, giving $\Pi^*(\tau) = \underline{\Pi}$ at every later charge.

For any policy optimum, baseline value, policy value, optimality and the baseline-output loss certificate give the successive inequalities

$$V(q_\tau, 0) \geq V(q_\tau, \tau) \geq V(q_0, \tau) \geq V(q_0, 0) - B_0(\tau)$$

Rearrangement proves (366). Substituting the quadratic separation premise and taking square roots proves (367). At and beyond a finite uniform cutoff, the entire objective and hence its complete argmax set are constant. Every contraction occurring later is consequently already represented at the cutoff in (46). Nonnegative output and the quantity bound yield the displayed upper bound on $\mathcal{D}_q$. No ordering of the maximizing quantities across different policy levels was used or follows from these arguments. □

Q.5.3 A resource bound on policy effects

The following result uses the implementation and contract assumptions of Online Appendix A.2.

Proposition Q.4 (A resource bound on policy effects)**.** *Normalize exit profit to zero. Let* $\mathcal{Q}$ *be the fixed set of attainable positive outputs and let* $\mathcal{T}$ *be a policy interval containing zero. Suppose the policy profit maximum is finite and attained for every* $\tau \in \mathcal{T}$*. Assume uniform constants* $\bar{v} > 0$ *and* $d > 0$*, and a policy-independent attainable continuation floor* F*, such that*

$$F(q) \leq W(q,\tau) \leq \bar{v}q, \qquad \mathcal{C}(q) \geq d \quad (q \in \mathcal{Q},\ \tau \in \mathcal{T}) \tag{368}$$

Define the protected profit

$$\Pi_F = \max\left\{0, \sup_{q \in \mathcal{Q}}[F(q) - \mathcal{C}(q)]\right\} > 0 \tag{369}$$

Every policy optimum is active and satisfies

$$q_\tau \geq \underline{q}_F := \frac{d + \Pi_F}{\bar{v}} > 0 \tag{370}$$

Consequently, for any baseline optimum q_0 and any $\tau \in \mathcal{T}$,

$$\frac{(q_0 - q_\tau)_+}{q_0} \leq 1 - \frac{d + \Pi_F}{\bar{v} q_0} \tag{371}$$

The supremum defining Π_F need not be attained.

Proof. Taking suprema in $W(q,\tau) - \mathcal{C}(q) \geq F(q) - \mathcal{C}(q)$ and including exit gives optimized profit at least Π_F, without requiring the floor supremum to attain. The finite policy maximum also bounds Π_F. Its strict positivity excludes exit. Every optimum therefore satisfies $\Pi_F \leq W(q_\tau,\tau) - \mathcal{C}(q_\tau) \leq \bar{v} q_\tau - d$, proving (370). The same bound at zero policy gives $q_0 \geq \underline{q}_F$. Thus $(q_0 - q_\tau)_+ \leq q_0 - \underline{q}_F$; division by $q_0 > 0$ proves (371). □

Q.5.4 Global policy ceiling

Proof. For each permitted capacity, Proposition Q.2 gives $W_n(q,\tau) \geq \underline{W}_n(q)$. Maximizing both sides over the same capacity menu proves (42).

Take the limit as $\tau \to \infty$ in (32). Proposition Q.2 identifies the limiting value as $\underline{W}_n(q)$. Since expected exposure is nonnegative, its improper integral exists and equals the finite receipt difference. It is positive exactly when $\bar{\chi}_n(q,t) > 0$ on a set of positive Lebesgue measure. Expected exposure records regulated orders in buyer-exclusion payment comparisons, including orders that lose in the actual allocation.

For any policy-optimal quantity q_τ, the remaining-market floor, optimality and weakly decreasing receipts give

$$V(q_\tau, 0) \geq V(q_\tau,\tau) = \Pi^*(\tau) \geq V(q_0,\tau) \geq V(q_0,0) - D_0$$

Exit is included with its unchanged payoff. These inequalities give the profit-loss and baseline-gap bounds. If $D_0 = 0$, baseline uniqueness forces $q_\tau = q_0$, and profit is unchanged. If every quantity outside the ε-neighborhood has gap greater than D_0, the gap bound excludes each such quantity.

Finally, if one baseline-optimal offered capacity n_0 at q_0 has zero cumulative exposure, then $\underline{W}_{n_0}(q_0) = W_{n_0}(q_0,0) = W(q_0,0)$. The optimized floor lies between this value and $W(q_0,0)$, so $D_0 = 0$. Thus positive exposure in every baseline-optimal selling choice is necessary for a production response. The protected marginal-return cases in Corollary A.21 show why it is not sufficient.

For the positive production floor, strict positivity of $\underline{\Pi}$ excludes exit. At any chosen active output,

$$\underline{\Pi} \leq W(q_\tau,\tau) - \mathcal{C}(q_\tau) \leq \bar{v} q_\tau - d$$

Rearranging gives (43). This is the resource bound of Proposition Q.4 with $F = \underline{W}$. It is uniform over the full tariff path and also applies to the auction with all regulated buyers removed. The guaranteed floor need not itself be feasible or attained; restrictions on feasible quantities can raise the smallest possible chosen output. □

Q.6 Physical and incentive architecture

Q.6.1 Conservation and retrospective exclusion

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 (45). 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. □

Q.6.2 Investigation and voluntary disclosure

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 (47). 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 (48). Its strict concavity establishes uniqueness; its first derivative gives (49). Implicit differentiation gives (69). 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. □

Q.6.3 Funded capture with optional suppression

Proof. For a fixed $\mathcal{T}$, original recipient i gains $\lambda_i t_i - a_i(\mathcal{T})$ relative to honesty. Acceptance therefore requires $t_i \geq a_i(\mathcal{T})/\lambda_i$. Each purchased reporter analogously requires $t_j \geq \sigma_j(\mathcal{T})/\lambda_j$. The joint implementability assumption makes these bounds sufficient as well as necessary. Thus every accepted schedule implementing $\mathcal{T}$ costs at least $\mathcal{B}(\mathcal{T})$ before settlement. If that bill exceeds F_0, no such schedule is funded. Otherwise equality payments

are funded and leave organizer surplus $G_\Omega(a) - a_0(\mathcal{T}) - \mathcal{B}(\mathcal{T})$; 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. □

Q.6.4 Positive-trade interior equilibrium

Proof. 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.

Strict convexity in (56) gives the buyer's unique best response. For an interior solution the uniform cost distribution gives $\bar{c}\pi = \Gamma - rD - D(1-r)a\pi$, yielding (58). The solution lies strictly inside the unit interval because its numerator is positive and less than $\bar{c}$; the last inequality in (57) ensures $e^* = a\pi^* < 1$. More generally, the fraud best response is the continuous, nonincreasing map $\pi \mapsto [\{\Gamma - DP(\min(a\pi, 1))\}/\bar{c}]_0^1$, 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 $p\sum_i b_i \le F$. The opening bargaining comparison establishes sequential reporting and the absence of mutually acceptable silence contracts. Posted wages, prepaid fees and segregated deposits satisfy the stated participation and budget contracts; buyer participation completes the positive-trade allocation.

For the matched benchmark, $\pi_0 = \widetilde{\Gamma}/\bar{c}$ and $e_0 = 0$. Equation (58) gives $\pi^* < \pi_0$ and $P^* > r$. The live-contamination function in (60) has derivatives $(1-P)/(1-\pi P)^2 > 0$ 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. □

Q.7 The certification price shock

Q.7.1 The faithful North–South price comparison

I maintain the faithful CES specialization of Section L.4: $B_C, B_D, I_N, I_S, \eta > 0$, $\varepsilon_{NS} > 1$, $0 \le \varphi_{NS} < 1$, and $\omega_{NS} > 0$. Set $I = I_N + I_S$, $\nu = I_N/I$, $a_C = (1-\varphi_{NS})^\eta B_C$, $a_D = B_D$, and $z_0 = (1-\varphi_{NS})^{-\eta/(1+\eta)}$. Then the dirty price falls relative to pooling exactly under

(206); equality leaves it unchanged and reversal raises it. Under a strict fall, the gross clean price satisfies (207). Regional output, receipts and profit satisfy (208).

Proof. For $z > 0$, the derivative of the clearing function is

$$\Upsilon'_{NS}(z) = (1-\nu)\frac{a_C}{a_D}(1+\eta)z^{\eta} + \omega_{NS}(\varepsilon_{NS}-1)z^{-\varepsilon_{NS}} > 0$$

The function tends to minus infinity at zero and plus infinity at infinity, so continuity gives exactly one positive root z^*. The pooling and separated dirty-price equations give

$$p_D < p_0 \quad \Longleftrightarrow \quad a_D + a_C(z^*)^{1+\eta} > B_C + B_D \quad \Longleftrightarrow \quad z^* > z_0$$

Strict monotonicity makes the last inequality equivalent to $\Upsilon_{NS}(z_0) < 0$. Substitution yields

$$\Upsilon_{NS}(z_0) = (1-\nu)\frac{B_C}{B_D} - \nu - \omega_{NS}(1-\varphi_{NS})^{\eta(\varepsilon_{NS}-1)/(1+\eta)}$$

which proves the stated three-way comparison. Furthermore

$$(p_C/p_0)^{1+\eta} = \frac{z^{1+\eta}(B_C+B_D)}{B_D + a_C z^{1+\eta}}$$

is strictly increasing in z. At $z = z_0$ it equals $z_0^{1+\eta}$, proving $p_C/p_0 > z_0$ when $z^* > z_0$. Finally the regional supply and optimized-profit formulas are proportional to p_D^{η} and $p_D^{1+\eta}$, respectively. Gross receipts have the same latter exponent. Taking ratios proves all three comparisons. The producer receives $(1-\varphi_{NS})p_C$, so the gross clean-price increase alone does not establish a higher net clean producer quote. □

Q.8 Research support

Q.8.1 Paid verification and clean-directed mineral research

The following result uses the implementation and contract assumptions of Online Appendix N.4.

Theorem Q.5 (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 $h_i(s) = 1 - e^{-\nu_i^R s}$, 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*

$$S_V + \sum_i N_i s_{i,t} \leq S \tag{372}$$

Choose a clean core $*$ and put $s_* = (S - S_V)/N_* > 0$. If

$$h'_*(s_*)[\underline{J}_* + G_*] - R'_*(s_*) > \max\left\{0, \max_{i \neq *}\left[h'_i(0)[\overline{J}_i + G_i]_+ - R'_i(0)\right]\right\} \tag{373}$$

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 (373); 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 N.18.

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 N.17. 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 $\epsilon > 0$. □

Q.8.2 Actual-value and uniform research-support thresholds

Use the complete owner-specific values in (221), evaluated at prices and a continuation consistent with the proposed full-clean allocation. Let $s_* > 0$, $h'_*(s_*) > 0$, differentiable convex laboratory costs, and concave increasing success probabilities satisfy the stated research assumptions. With $\mathcal{B}_t$ defined in (226), the allocation is weakly supported exactly when $G_{*,t} \geq \max\{0, \Gamma_t\}$, subject to $G_{*,t} \geq 0$. A strict marginal advantage requires $G_{*,t} > \Gamma_t$. The bound (227) supplies the stated uniform sufficient condition.

Proof. The clean owner can employ s_* at a supporting scientific wage if and only if its marginal return at s_* is at least every rival's return at zero effort and the zero return to unused labor. This is

$$h'_*(s_*)[\widehat{J}_{*,t} + G_{*,t}] - R'_*(s_*) \geq \mathcal{B}_t$$

For a nonnegative reward $\widehat{J}_i + G_i$, the owner's objective is concave, so these marginal conditions are sufficient for its optimum. A rival with negative reward has a nonincreasing objective because success probability and research costs are nondecreasing; zero effort is optimal. The comparison with zero in $\mathcal{B}_t$ covers that case. The stated scientist accounting makes the allocation feasible. Rearranging and dividing by $h'_*(s_*) > 0$ gives $G_{*,t} \geq \Gamma_t$. Intersecting with nonnegative payments yields the weak minimum, and the same algebra with a strict inequality yields the strict condition. Therefore $\Gamma_t < 0$ permits a strict advantage at zero clean support, $\Gamma_t = 0$ gives equality there, and $\Gamma_t > 0$ requires positive support for this proposed allocation. Replacing the clean value by a valid lower bound and competing values by valid upper bounds yields the uniform sufficient comparison. Failure of that bound does not establish failure at actual equilibrium values. The calculation must be repeated if policy changes those values. □

Q.9 Path dependence and private takeover

Q.9.1 Capacity, private takeover and the duration of support

I maintain the fixed-budget faithful CES and one-interval ownership cancellation assumptions of Section L.11. They are not asserted for the unrestricted VCG or general-contract economy. With its constants $u_C, u_D > 0$, $c_C, c_D \geq 0$, inherited relative capacity $a_t > 0$, and full-clean growth factor $\bar{v}_C > 1$, the expenditure and capacity cutoffs are (231) and (232). The private, temporary-support and continuing-support cases stated there follow.

Proof. With unit mass of owners and full-clean aggregate factor $\bar{v}_C$, the unsuccessful clean owner's operating value is $(1 - \varphi_{NS})E/[(1+\eta)\bar{v}_C]$. Its owned productive improvement adds γ_C times that value. The dirty aggregate factor is one, giving owned increment $\gamma_D(I-E)/(1+\eta)$. The stated one-interval cancellation conditions remove later differential ownership and contract values; subtracting the dated installation bill K_j and discounting gives (229). In particular, the unsuccessful clean owner's capacity accounts for the divisor $\bar{v}_C$. Substitution into the marginal research conditions gives the clean return $u_C E - c_C$ and competing dirty return $u_D(I-E) - c_D$. The two strict inequalities $u_C E - c_C > 0$ and $u_C E - c_C > u_D(I-E) - c_D$ are equivalent to

$$E > \max\{c_C/u_C, (u_D I + c_C - c_D)/(u_C + u_D)\} = E^\dagger$$

CES expenditure shares and clearing imply

$$a(E) = \frac{E}{I-E}\left[\frac{E - I_N}{\omega_{NS}(I-E)}\right]^{(1+\eta)/(\varepsilon_{NS}-1)}, \qquad I_N < E < I$$

Both positive ratios in this expression increase strictly with E. The map is continuous, tends to zero as $E \downarrow I_N$, and tends to infinity as $E \uparrow I$. Hence it is invertible. Research at t earns its production return at relative capacity $\bar{v}_C a_t$, so $E > E^\dagger$ is equivalent, for an interior cutoff, to $a_t > a(E^\dagger)/\bar{v}_C = a^\dagger/\bar{v}_C$.

Along a funded full-clean path, repeated application of (228) gives $a_t = \bar{v}_C^t a_0$. If $I_N < E^\dagger < I$, geometric growth crosses the finite cutoff at the first T such that $\bar{v}_C^{T+1} a_0 > a^\dagger$. Thereafter expenditure increases, strengthening the strict private comparison. If $E^\dagger \le I_N$, the private comparison is already strict at every finite positive capacity state.

As capacity diverges, $E \uparrow I$ and the dirty marginal return tends to $-c_D \le 0$. The clean marginal return tends to $u_C I - c_C$. Thus finite strict private takeover is possible exactly when $u_C I - c_C > 0$. At equality, the clean return remains below zero at every finite state but tends to zero, so the required additive support remains positive and tends to zero. If $u_C I - c_C < 0$, its limiting requirement is positive. Adequate funding can still sustain the selected allocation. These statements concern existence and persistence of that continuation, not uniqueness, disappearance of dirty production, or physical conservation. □

Q.10 The government intervention threshold

Q.10.1 Eventually strict intervention on the recleared path

The following result uses the implementation and contract assumptions of Online Appendix N.4.

Theorem Q.6 (Eventually strict intervention on the recleared path)**.** *In any connected equilibrium suppose each type has a funded legal-flow floor $\widehat{R}_{a,t} \ge \underline{R}_a > 0$ and normalized protection flows satisfy $0 \le \widehat{b}_{a,t} \to 0$ uniformly over feasible population paths. If*

$$k_a(1 - \widehat{\delta}_G) < \widehat{\delta}_G P_a \underline{R}_a \tag{374}$$

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 $D_{a,t} = U_{a,t} - V_{a,t}$:

$$D_{a,t} = \min\{\widehat{R}_{a,t} - \widehat{b}_{a,t} + \widehat{\delta}_G D_{a,t+1}, \widehat{R}_{a,t} + k_a + \widehat{\delta}_G(1 - P_a) D_{a,t+1}\}$$

If $\widehat{b}_{a,t} \le b_*$ for all $t \ge T$, comparison of this increasing contraction with its constant-flow version gives

$$D_{a,t} \ge D_* := \min\left\{\frac{\underline{R}_a - b_*}{1 - \widehat{\delta}_G}, \frac{\underline{R}_a + k_a}{1 - \widehat{\delta}_G(1 - P_a)}\right\}, \qquad t \ge T$$

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 $-k_a - \widehat{b}_{a,t} + \widehat{\delta}_G P_a D_{a,t+1}$,

at least $-k_a - b_* + \widehat{\delta}_G P_a D_*$. At $b_* = 0$ this lower bound is strictly positive exactly when (374) 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 $(1 - P_a)^n \to 0$. □

Q.10.2 The fiscal threshold with waiting and repeated attempts

I maintain the intervention timing, funded fiscal flows and Bellman conditions of Section M.2. For constant normalized flows, write $d \in (0, 1)$ for the continuation factor, $P \in (0, 1)$ for success probability and $k > 0$ for the attempt cost. The stationary threshold is (241), with tax-base form (242). On the changing market path, (245) is a sufficient bound, with the true choice determined by (239).

Proof. Constant legal flow gives $U = R/(1 - d)$. Waiting forever gives $V^{\text{wait}} = b/(1 - d)$, and attempting after each failure solves $V^{\text{try}} = -k + d[PU + (1 - P)V^{\text{try}}]$, so

$$V^{\text{try}} = \frac{-k + dPR/(1 - d)}{1 - d + dP}$$

Both denominators are positive. Multiplying the comparison with V^{wait} by them and collecting terms yields $dP(R - b) \geq (1 - d)(k + b)$. With $R = \xi_a \widehat{\pi}^C - f$ and $\xi_a > 0$, solving for $\widehat{\pi}^C$ gives the printed tax-base threshold. Equality is indifference and a strict inequality gives strict intervention.

For varying flows, subtraction of the Bellman equations gives (243). Under the funded lower bound $\underline{R}_a$ and tail upper bound b_*, the monotone discounted comparison gives (244). For its waiting branch, the intervention advantage is strictly positive precisely when

$$dP\underline{R}_a - (1 - d)k_a - (1 - d + dP)b_* > 0$$

For its intervention branch the same numerator appears after multiplication by the positive $1 - d + dP$. Thus imposing positivity of this common numerator suffices for either branch, and solving for b_* proves (245). The legal-floor crossing condition makes the critical value positive; uniform convergence of protection income to zero then gives a finite tail on which strict intervention holds. This does not replace the anticipated changing-path continuation values by a stationary annuity before that tail. □

Q.11 Conservation

Q.11.1 What governance does and does not conserve

The following result uses the implementation and contract assumptions of Online Appendix N.4.

Proposition Q.7 (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,

$$\sum_a \bar{d}_a M_{a,0}^D \left(T + \frac{1}{P_a} \right) \leq \mathcal{B} \tag{375}$$

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. □

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, Liguo 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.

Shawn Davies, Therése Pettersson, and Magnus Öberg. Organized violence 1989–2025, and violent political protests. *Journal of Peace Research*, 63(4):705–723, 2026. doi: 10.1093/jopres/xjag046.

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 DRCongo'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 Esquível-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.

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, 2021a. URL https://doi.org/10.1787/4d432567-en.

OECD. Gold flows from Venezuela: Supporting due diligence on the production and trade of gold in Venezuela. OECD Publishing, Paris, 2021b. URL https://doi.org/10.1787/7a392f42-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.maapprogram.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.

Cédric Perrin, Étienne Blanc, Olivier Cigolotti, and André Guiol. Le Plateau des Guyanes : affirmer la France comme un acteur régional clé. Rapport d'information 135 (2025–2026), Sénat, 19 November 2025. URL https://www.senat.fr/rap/r25-135/r25-1351.pdf.

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.

Ahmed Soliman and Suliman Baldo. Gold and the war in Sudan: How regional solutions can support an end to conflict. Research paper, Chatham House, London, March 2025. URL https://www.chathamhouse.org/2025/03/gold-and-war-sudan.

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.

Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992. https://dash.harvard.edu/entities/publication/73120378-cb4d-6bd4-e053-0100007fdf3b.

Benny Kleinman. Wage inequality and the spatial expansion of firms. Working paper, September 2025. https://binyamink.github.io/bennykleinman/bk_JMP.pdf.

Klaus Desmet and Esteban Rossi-Hansberg. Spatial development. *American Economic Review*, 104(4):1211–1243, 2014. https://rossihansberg.economics.uchicago.edu/SD.pdf.