MONTREAL.AI · QUEBEC.AI

INSTRUMENT EDITION
1.1

SUCCESSOR Ω

DAEDALON Ω

THE CAPACITY FOUNDRY

Executable understanding. Qualified output. Productive renewal.

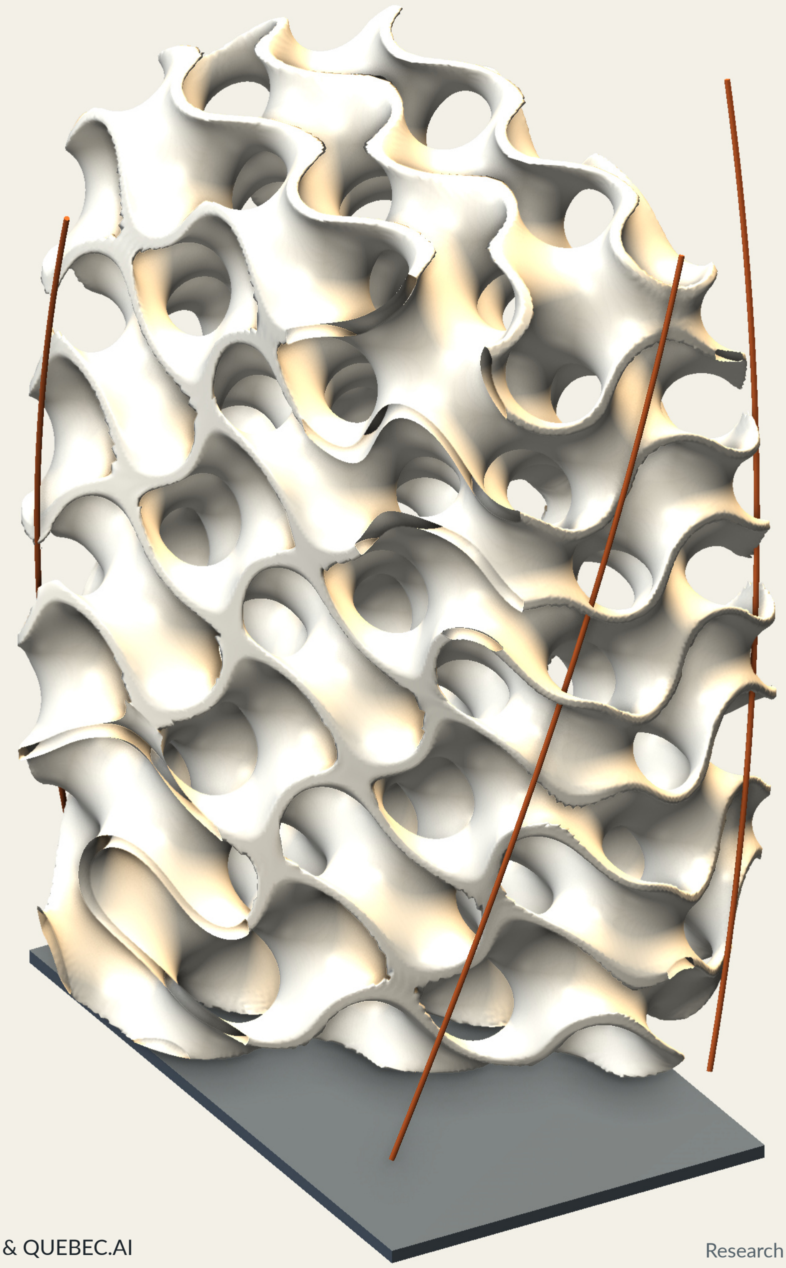

Vincent Boucher
President, MONTREAL.AI & QUEBEC.AI

DAEDALON-TR-01
Research 1.0 / visual revision 1.1

Conceptual parametric geometry. Not a discovered material or a built instrument.

Publication record and evidence contract

Title	SUCCESSOR Ω — DAEDALON Ω
Subtitle	The Capacity Foundry: From executable process models to a renewable productive estate
Author	Vincent Boucher, President, MONTREAL.AI & QUEBEC.AI
Research edition	DAEDALON-TR-01 / 1.0 / 5 September 2026
Visual revision	Instrument Edition 1.1; unchanged research record
Implementation	NEURAL-SYMBOLIC SUCCESSOR Ω 10.2.0
Empirical basis	DAEDALON simulation 1.0; eight authored synthetic cohorts
New contribution	Conditional mathematics and executable inspection/renewal studies
Status	AI-assisted research manuscript; not independently peer reviewed

The proposition

Own the capacity to learn which physical process deserves to become a productive asset, then retain the knowledge, evidence and replacement capability that keep that asset useful.

This paper develops an industrial research argument from a delivered executable study. It preserves its negative competitive result and its limited operating domain. The added mathematical statements are conditional derivations and specialized consequences of established methods, not assertions of previously unknown mathematics. The new computations do not change the first-party product.

Evidence class	Meaning in this publication
Executed	Local fitting, neural training, symbolic revision, decision calculations, tests and fresh replay.
Derived	Statements proved under explicit mathematical assumptions; numerical consistency checks.
Conditional	Authored process laws, prices, 24-year cash scenario, inspection prior, defect costs and resource envelopes.
Prospective	Real customer proof, independently admitted Specialist ASI, factory control, autonomous assembly and stellar infrastructure.

No source of funding, customer acceptance, apparatus deployment, author signature, independent attestation or economic receipt is invented. Product and third-party licences remain applicable. Exporting a ZIP does not itself establish legal ownership. Editorial images combine new parametric geometry and an image-only AI-generated fabrication crop; they are not installations or technical evidence. Every numerical plot is computed separately from disclosed data or equations.

Abstract

SUCCESSOR Ω—DAEDALON Ω examines a customer-controlled institution that learns manufacturing response functions, converts them into objective-specific decisions, and accounts for the output, cash and replacement flows required to maintain an expanding productive estate. The retained study uses the unchanged scientific engine of NEURAL-SYMBOLIC SUCCESSOR Ω v10.2. Eight cohorts train eight local neural guides and three symbolic generations each. Neural ablation changes the proposed observation in every cohort. New observations falsify the initial hypotheses; an authored regime change subsequently falsifies every second-generation model. Third-generation candidates recover local predictive validity without inheriting operational authority.

The primary fulfillment comparison improves authored annual costs by CA$1.167 billion against a fixed-setpoint incumbent, but loses CA$4.326 million after incremental costs against selected generic Beta and CA$7 million against a mechanism-informed conventional comparator. The eight-cohort bootstrap interval for the generic comparison is [CA$-8.022, CA$-0.987] million. A separately constituted own-sales mission increases mean accepted output by 4.15% and reduces electricity per line-week by 22.58%, producing a CA$599.587 million incumbent contrast at the assumed scale. Informed Beta matches the nominal physical choices. These alternative contrasts are not additive and do not establish proprietary Alpha.

New publication-side calculations distinguish measured acceptance from latent quality. A beta-binomial example raises the probability of missing 1,000 good units from 5.03% to 30.54% when intralot correlation rises from zero to 0.01, while preserving mean output. An imperfect-inspection decision study selects 200 extra coupons under a 1% posterior release-risk gate, with modeled expected cost CA$585,914 versus CA$950,000 for automatic rework. Its nominal joint unsafe-release probability is 0.147%; deliberate prior and instrument misspecification raises that quantity to 1.849%. Those conditional calculations are not a qualified sampling plan.

The formal treatment derives capped-sales and decision-error bounds, exergy dissipation, thermal barrier conditions, inspection incentives, information-acquisition rewards, congestion potentials, finite Gibbs coordination, a successor-wide error budget and age-structured productive renewal. A retained 24-year cash scenario produces four additional operating factories relative to fixed-setpoint operation, but less closing cash after greater investment; informed Beta remains cheaper. A stellar resource envelope explicitly separates useful power, captured power, heat rejection, material stock and replacement throughput. Its horizon is a research programme, not an arrival forecast. No independent Specialist ASI, unrestricted theory discovery, field Alpha or production authorization is asserted.

Keywords. Neural-symbolic learning; manufacturing; imperfect metrology; beta-binomial quality; inspection games; mission Alpha; delayed capacity; replacement throughput.

Owner decision brief

The next asset is the ability to replace the wrong answer

Retain the trained knowledge and the strongest available supplier. Fund a new mission only when its expected decision value, measurement quality and feasible operating controls justify the burden. Do not pay an exclusive-model premium that the comparative evidence rejects.

Decision-relevant result	Value
Primary mean versus generic Beta, after annual costs	CA$ -4.326 M
Primary mean versus informed Beta, after annual costs	CA$ -7.000 M
Own-sales contrast versus fixed setpoint, assumed annual scale	CA$ 599.587 M
Operating factories after 24 scenario years: candidate / fixed	22 / 18
Candidate cash difference versus fixed at that endpoint	CA$ -2.722 B
Actual customer Alpha credited	CA$ 0
Independent admission / production authority	None

The valuable next mission

Measure the distinction between latent defects and the inspection label; establish cross-unit and cross-site dependence; test whether better metrology or a physically different process changes the available decision set. The new inspection laboratory gives a reproducible example of that question, not a ready-made factory release policy.

Three reading routes

Owner: Sections 1, 6, 8, 18–24 and the six-page decision book. **Scientist**: Sections 3–12 and the derivations. **Engineer and verifier**: the source map, protocol, adverse cases, exact file commitments and reproduction instructions. Every path ends at a named decision, not a generic claim of intelligence.

READER'S INSTRUMENT / THREE SCALES

An image opens the question. The evidence carries the claim.

New AI-generated fabrication concept; image-only crop. No depicted apparatus was built or tested.

01 / Manufacture

Sections 1–6: inspect the three generations of world programs, the changed observations and the two distinct economic objectives. The trained candidate remains a research object; the informed Beta comparison remains negative after its incremental cost.

02 / Measure and decide

Sections 7–17: distinguish a latent defect from an inspection label; inspect the conditional release-risk study, energy balances, quality incentives and authority boundaries. Mathematical validity is not independent verification of the assumed world.

03 / Pay and renew

Sections 18–24: follow the paid construction account, age-structured stock and replacement requirements. Conditional capital paths and stellar arithmetic are not receipts, installations or forecasts.

Revision boundary. Edition 1.1 rebuilds the visual and navigational layer. It does not change v10.2, the retained experimental observations, numerical results, 14 propositions or qualification status. The mathematical studies and the manufacturing source remain separately inspectable.

Contents

1 The institution that learns what to manufacture

OWNER SCENARIO / FICTIONAL CONTEXT, EXECUTED RESEARCH CORE

Dr. Mira Sorel, a fictional physicist and manufacturing-control expert, directs an industrial group with CA$320 billion in assets, 32 factories, eight lines per factory and CA$14 billion in liquidity. These are setting assumptions, not financial facts about the author or user. Her problem is not whether capital can purchase a machine. It is whether the organization understands a changing process well enough to decide which machine, which setpoint and which production commitment deserve capital.

Mission charter

Learn what changes the operating decision. Count only accepted physical output. Compare with the strongest feasible alternative. Never spend future revenue.

The customer initially asks for a lower-cost way to fulfill a delivery obligation. Subsequently, she constitutes a different mission: earn surplus from a factory's own accepted products. The distinction is economically material. The first mission can buy replacements when its own production is insufficient. The second must not call those purchases manufactured output or credit them to a capacity-reproduction story. The retained development record records this objective change rather than transferring an attractive result to an incompatible business.[1]

Her durable asset is the constituted mission and its supporting programs, observations, trained parameters, alternatives, falsifiers and failure history. The source architecture separates this customer-owned layer from replaceable cognition and from accountable authority. Its page-12 architecture and page-21 serving/learning diagram make the boundary explicit: the branch that learns cannot silently modify the branch permitted to act.[2]

The productive thesis is deliberately larger than one predictor: a customer may compound methods, measurement quality, production know-how and recovery capabilities while replacing the computational supplier. It is also narrower than a promise of automatic wealth. Each claim of additional value needs its own counterfactual, time unit, full cost, physical feasibility and admission envelope.

2 What the retained software actually executes

The experiment uses the scientific branch of SUCCESSOR Ω v10.2, not an independently deployed factory controller. All 55 first-party payload files in the delivered source snapshot are retained unchanged. The application code is a separate customer extension: it supplies the synthetic process, finite decisions, monetary accounting, second-mission objective and capital scenario. No apparatus, procurement, settlement or robotic connector is invoked.[3, 1]

Artifact or module	Role and limitation
`omega.science.search`	Typed expression search, coefficient fitting, evaluation and lineage; finite language and budget.
`omega.science.neural`	Locally trained surrogate and experiment proposal; not a frontier-model service.
`src/manufacturing.py`	Authored plant, fulfillment objective and separate ledger reconstruction.
`src/market_mission.py`	Exploratory own-sales objective, frozen models and new cases.
`src/capacity.py`	Conditional local investment screen; no global lifetime optimum.
`research/qualification_frontier.py`	New publication-side inspection and renewal calculations; not native product controls.

Each of eight cohorts uses a different declared synthetic coefficient tuple. Data identities distinguish development from nominal holdout, stress and shifted-regime observations. The complete archive remains in the same author's custody. A procedural freeze and disjoint identifiers are useful engineering controls, but they do not establish independent sequestration, an externally pre-registered test or real industrial causality.

For this publication the complete primary study was rerun, reproducing 155 output files byte-for-byte. The second mission and capital continuation reproduced a further ten files. The 68 extension tests were rerun. These are fresh same-environment reproducibility results, not a new independent confirmation cohort. Historical upstream counts are not relabeled as newly executed tests.

3 Three generations of executable understanding

The synthetic acceptance law is deliberately expressible by the supplied language. For normalized process intensity $u \in [0.2, 2.8]$ and observed contamination $z \in [0, 1]$, cohort c has a nominal mean response

$$p_c(u, z) = a_c + b_c \sin u + c_c z + d_c uz, \qquad p_c^{\text{shift}}(u, z) = p_c(u, z) - 0.12z^2. \tag{3.1}$$

The coefficients differ modestly across cohorts. Coupon means include uniform instrument noise bounded by 0.001. This observation noise is distinct from binomial production noise. Neither the normalized variable nor the equation is an operating recipe for a real material.

G1 uses 32 training and 20 selection observations from a narrow domain. In cohort zero it selects

$$\hat{p}_1(u) = 0.69291895 + 0.17595799u. \tag{3.2}$$

A 16-hidden-unit tanh neural surrogate is trained for 1,600 steps; it contains 65 fitted weights and biases, plus normalization state. Its disagreement with G1 contributes to choosing one of 160 permitted probes. In cohort zero the probe is approximately $(u, z) = (2.78719, 0.68244)$ and contradicts G1 by 0.55428 in accepted fraction. Removing the neural contribution changes the proposed probe in all eight cohorts. That ablation establishes computational influence, not optimal experimental design or causal attribution of financial gain.

G2 receives broader training observations and the probe. The cohort-zero selected model is

$$\hat{p}_2(u, z) = 0.67957861 - 0.054964397z - 0.046805831zu + 0.21929037 \sin u. \tag{3.3}$$

All eight G2 candidates pass the declared nominal equation protocol. When the separate process shift is introduced, all eight fail it. G3 is fitted on shifted development observations and challenged on new shifted cases. Every G3 passes its local protocol; none inherits authority. One G3 represents part of the quadratic shift by a cosine approximation. A finite predictive pass therefore does not identify a unique mechanism.

Sparse equation discovery and neural-assisted symbolic regression provide related technical precedents; code-world-model research explicitly uses executable representations in planning. DAEDALON uses its own restricted implementation and synthetic adapter, rather than claiming that those publications prove this complete institution.[4, 5, 6]

4 Stochastic manufacturing and conserved accounts

The production catalogue contains 27 intensities $u = 0.2, 0.3, \ldots, 2.8$. For a line-week, the authored attempt count, electrical consumption and feedstock are

$$n(u) = \left\lfloor \frac{1800}{1 + 0.28u} \right\rfloor, \qquad E(u) = n(u)(10 + 12u^2)\ \text{kWh}, \tag{4.1}$$

$$M_{\text{input}} = 35n(u)\ \text{kg}, \qquad G \mid u, z \sim \text{Binomial}(n(u), p_c(u, z)). \tag{4.2}$$

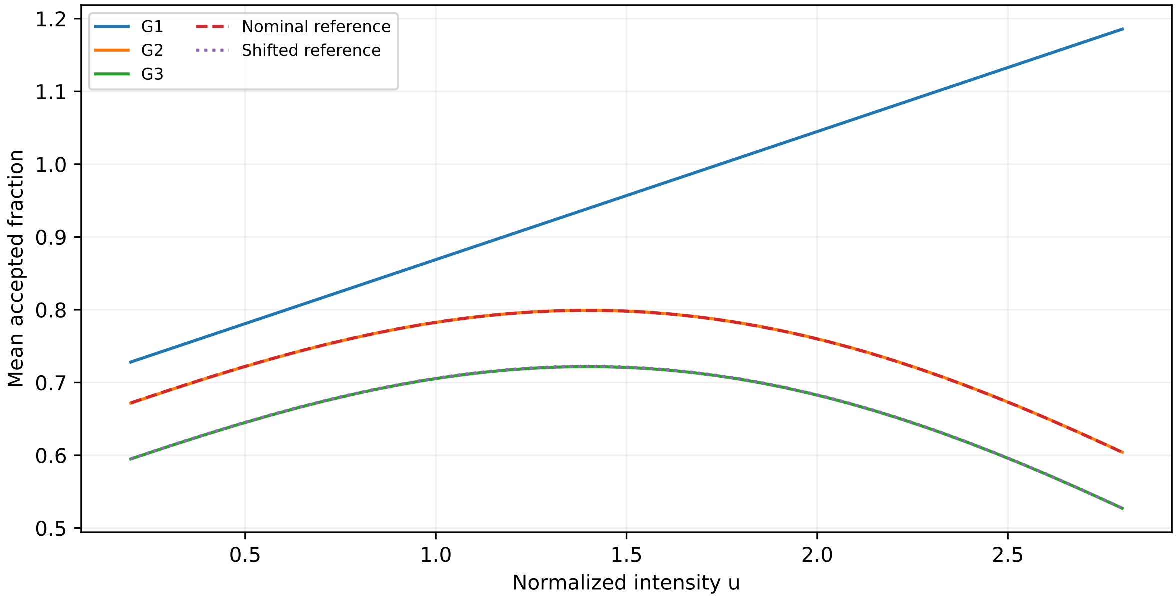

Figure 1. Frozen hypotheses and the authored process at contamination 0.8, cohort zero. G1 is displayed without clipping away its failure.

Every rejected unit remains paid input. Accepted mass is $35G$ and scrap mass is $35(n-G)$. A separate ledger reconstructs these quantities, electrical expense, raw-material cost, quadratic wear cost, staffing, purchases and unmet-demand charges. A shared exogenous uniform couples paired policy realizations. This supports low-noise comparisons under the authored stochastic model; it is not multiple independent factory observations.

For demand D and replacement capacity K, write $O = \min(G, D)$, $B = \min(K, D-O)$ and $U = D-O-B$. The fulfillment cost is

$$C_{\text{ful}} = nc_m + 45nu^2 + \pi_E E + 110000 + c_b B + (c_b + 25000)U. \tag{4.3}$$

Monetary coefficients are synthetic CAD. Surplus and scrap have zero sale credit. Replacement purchases are paid and cannot later be counted as own production. The nominal replacement capacity is 800 units; the supply-shock scenario reduces it to 120. The model omits queues, correlated lot defects, measurement misclassification and transport timing; Sections 7–9 isolate some consequences of those omissions without silently rewriting the original study.

Accounting invariant

One physical unit may occupy one disposition at a time: delivered externally, retained internally, scrapped or held. A receipt is not a second unit. A capacity promise is not an installed machine. A future sale is not current cash.

5 The objective is part of the world model

A factory selling only its own panels faces a different objective. For a unit price r and market limit D_s, it selects

$$u^* \in \arg\max_{u \in \mathcal{U}} \{r\mathbb{E}[\min(G, D_s)] - nc_m - 45nu^2 - \pi_E E - 110000\}. \tag{5.1}$$

The exploratory market mission sets $r = \text{CA\$}\,3300$ and $D_s = 1600$. It uses frozen G2 models on 1,024 newly generated cases. Generic Beta is retained from the fulfillment selection rather than reselected for the new business, a limitation explicitly addressed by retaining the informed comparator. A different objective is a new mission, even when the predictive model is unchanged.

Proposition 5.1 (Exact capped output and marginal information value)**.** Let $G \sim \mathrm{Binomial}(n, p)$ and integer $1 \le d \le n$. With binomial CDF F,

$$v_d(p) = \mathbb{E}\min(G, d) = \sum_{k=1}^{d} \mathbb{P}(G \ge k) \tag{5.2}$$

$$= npF(d-2; n-1, p) + d[1 - F(d-1; n, p)], \tag{5.3}$$

$$v'_d(p) = nF(d-1; n-1, p), \qquad 0 \le v'_d(p) \le n. \tag{5.4}$$

The tail-sum identity counts each delivered unit once. The derivative shows why a uniform gain in prediction accuracy has nonuniform decision value: near a binding sales cap, another potential accepted unit may have almost no revenue value. Conversely, near a costly delivery threshold, the same prediction error can be consequential. The proof is in Appendix D.

Proposition 5.2 (From yield error to decision regret)**.** For a finite catalogue, assume $|p(u) - \hat{p}(u)| \le e$ uniformly, known fixed costs, and a payoff that is L-Lipschitz in the accepted count. If each action has at most $n_{\max}$ attempts, the action optimal under the approximate binomial model has true expected regret at most $2Ln_{\max}e$.

This is a conditional decision bound, not a guarantee derived from average test error. It requires uniform coverage of the actual decision catalogue and the binomial coupling. It does not survive unmodeled dependence merely because the means remain correct. Risk gates must still be checked separately.

6 Results: productive improvement is not exclusive Alpha

The eight primary cohorts execute 15,610 native candidate fits, 96 generic comparator fits and eight separately fitted mechanism-informed controls. Each cohort chooses its yield margin using the same 48 decision-selection cases. Nominal evaluation then uses 128 cases per cohort, with distinct shifted and supply-shock cases. The generic alternatives include polynomial ridge, ExtraTrees and histogram-gradient boosting. The informed conventional control receives the true nominal feature family $[1, \sin u, z, uz]$, but fits its coefficients on the same observations. Its additional knowledge is disclosed, not represented as blind discovery.

Primary fulfillment contrast	Before costs	After costs
Candidate versus fixed setpoint	CA$ 1167.417 M	Different baseline
Candidate versus selected generic Beta	CA$ 2.674 M	CA$ -4.326 M
Candidate versus informed Beta	CA$ 0	CA$ -7.000 M

Annualization multiplies a line-week mean by $32 \times 8 \times 52 = 13,312$, an authored scale convention. The CA$7 million recurring deduction consists of CA$4 million operations/human burden and CA$3 million reserves. A separate first-year CA$6 million research/integration allowance yields CA$-10.326 million versus generic Beta and CA$-13 million versus informed Beta. These are not field receipts.

Only one primary cohort wins against selected generic Beta after costs. The conditional bootstrap resamples eight cohort means 20,000 times; it does not treat 1,024 policy cases as 1,024 independent training experiments. The resulting interval excludes zero on the negative side. Since informed Beta matches the candidate's nominal physical choices, no premium can be justified by exclusive predictive performance in this trial.

The own-sales mission increases accepted own output from 1,091.925 to 1,137.289 units per line-week and lowers electricity from 39,939.32 to 30,919.10 kWh. Mean operating surplus rises by CA$45,041.10 per line-week relative to fixed operation. Informed Beta again matches the nominal candidate. The 17,408 combined

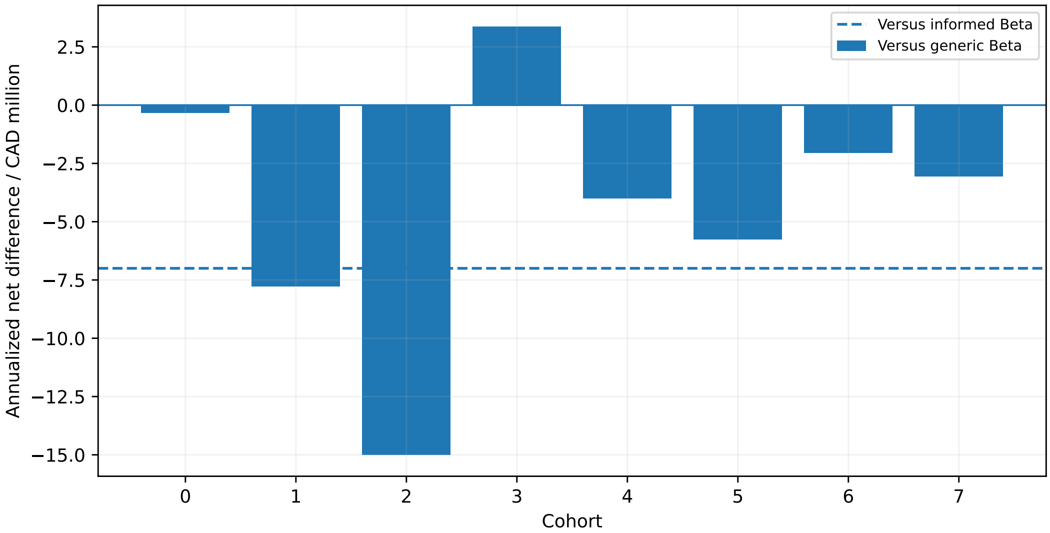

Figure 2. Eight cohort-level annualized comparisons after recurring costs. The zero line is the break-even threshold, not an admission certificate.

ledger checks and 1,536 detected corruptions establish tested accounting consistency, not physical certification. Under supply shock, 6,618 candidate units remain unmet. The protocol retains that failure.

Acceptance is a measurement.
Quality remains a claim.

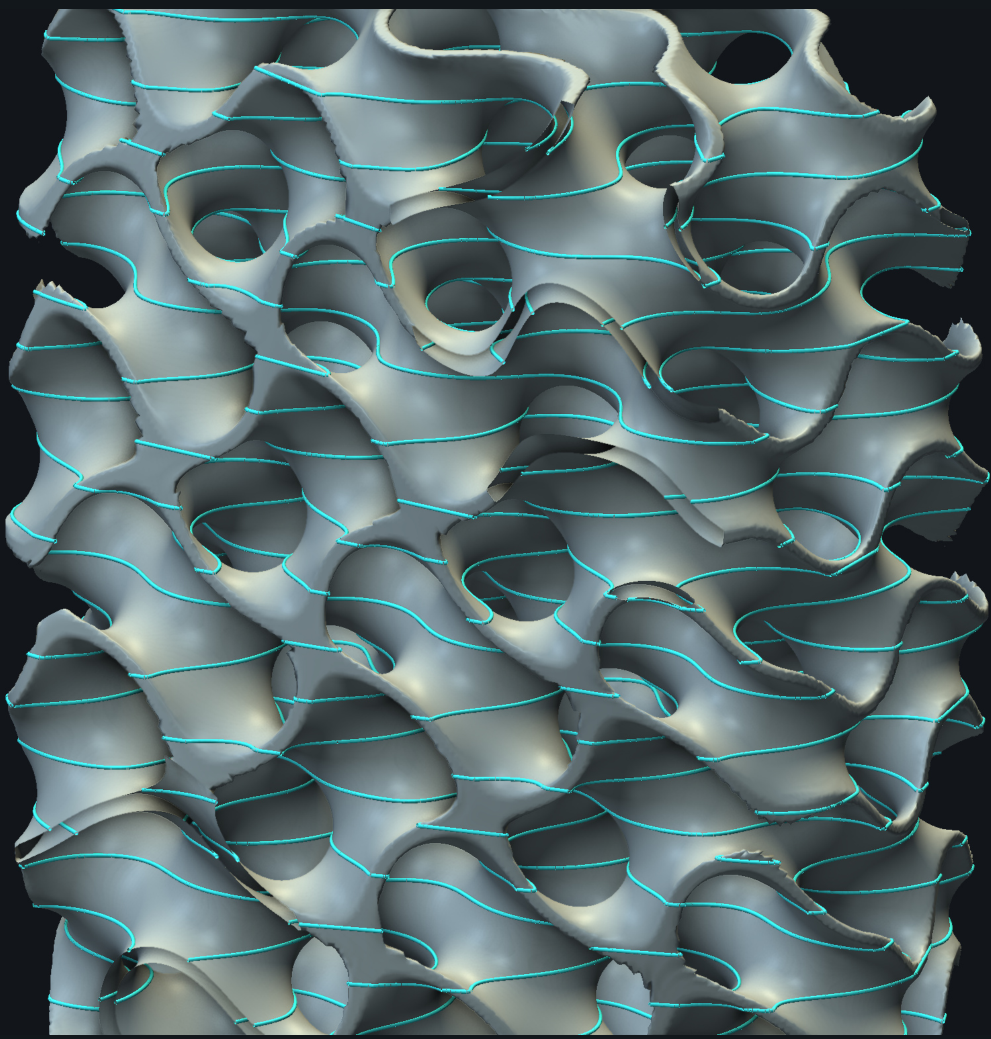

A state, a sensor and a release decision are different objects. The following chapters keep their uncertainties separate.

Conceptual geometry with computed surface-intersection contours. Not a scan, a test result or metrology evidence.

7 The metrology gap: accepted is not necessarily good

In the retained simulator, acceptance is a direct stochastic draw from the authored process. A real instrument instead produces a label. If the label has false negatives, drift or a shared calibration error, accepted count may overstate latent functional output. This is not corrected by a more elaborate economic optimizer. Acceptance sampling concerns a decision about a lot, not a claim that every accepted item is defect-free.[7, 8]

Proposition 7.1 (Quality conditional on an acceptance label)**.** Let true good fraction be p. An inspection rejects a defective item with sensitivity s and falsely rejects a good item with probability f. If $p(1-f)+(1-p)(1-s)>0$,

$$\mathbb{P}(\text{good} \mid \text{accepted}) = \frac{p(1-f)}{p(1-f)+(1-p)(1-s)}. \tag{7.1}$$

Bayes' rule proves the formula. It shows why increasing apparent acceptance by weakening the detector can reduce delivered quality. A target based only on the label creates an incentive to improve the measurement score instead of the object. Qualification therefore needs references that the production optimizer and inspector cannot redefine.

The next empirical mission would characterize sensitivity and specificity across operating conditions, estimate uncertainty in those quantities, validate reference specimens, and monitor shared calibration drift. The present extension does not obtain those measurements. It exposes their decision consequences using openly invented channel parameters.

8 Correlated lots and risk-limited release

A shared lot condition creates dependence even when items are independent conditional on that condition. Let latent good probability $P \sim \text{Beta}(a,b)$, and $G \mid P \sim \text{Binomial}(n,P)$. Write $p = a/(a+b)$ and $\rho = 1/(a+b+1)$.

Proposition 8.1 (Variance inflation at unchanged mean)**.**

$$\mathbb{E}G = np, \qquad \text{Var}(G) = np(1-p)[1+(n-1)\rho], \qquad n_{\text{eff}} = \frac{n}{1+(n-1)\rho}. \tag{8.1}$$

The last expression is the independent-Bernoulli sample size with the same variance of the sample mean, not a count of independently observed items.

For a new authored illustration with $n = 1200$, $p = 0.85$ and target 1,000 good units, the shortfall probability is 5.03% under independence. At $\rho = 0.01$ it is 30.54%, with the same expected count of 1,020. Standard deviation grows from 12.37 to 44.58 units, and n_{eff} falls to about 92.38. One thousand two hundred same-lot items are not a substitute for many independent lots.

8.1 A priced coupon-and-rework decision

The second new calculation uses defect probability q, not the original acceptance p_c. It assumes a 1,000-unit shipment, $\mathbb{E}q = 0.02$, $\rho = 0.01$, sensitivity 0.95 and false-alarm rate 0.005. The owner can purchase $m \in \{0, 10, 20, 40, 80, 120, 200\}$ extra same-process coupons at CA$180 each. These coupons are not removed from the shipment. Conditional on q, flagged coupon count K has

$$K \mid q \sim \text{Binomial}(m, h(q)), \qquad h(q) = 0.005 + (0.95 - 0.005)q. \tag{8.2}$$

After observing k, the finite policy either releases the shipment or pays CA$950 per unit for complete rework/replacement. A released defective unit incurs a CA$25,000 consequence. The toy rework branch is

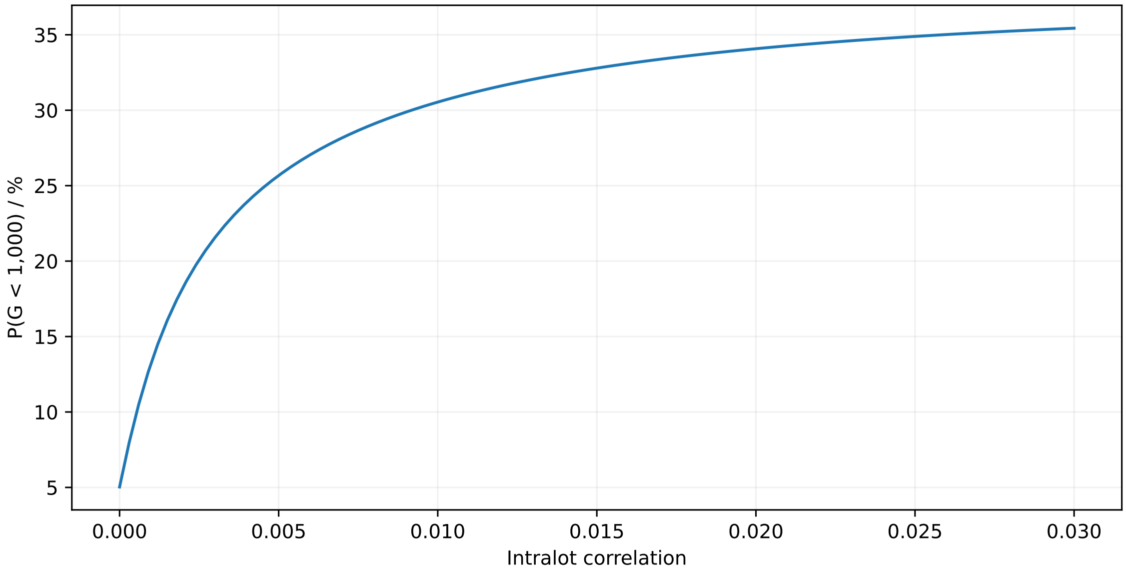

Figure 3. A small intralot correlation changes the delivery tail while preserving expected output. New authored beta-binomial calculation, not a re-estimate of the original process.

assumed perfect and fully available; qualifying real rework, delay and capacity is explicitly outside this calculation.

$$\mathcal{C}(m) = 180m + \sum_{k=0}^{m} \mathbb{P}(K = k) \min\{25000N\mathbb{E}[q \mid k],\ 950N\}, \qquad N = 1000. \tag{8.3}$$

Without a release-risk constraint, the cheapest menu option is no test and release, with CA$500,000 expected cost and 4.245% probability of more than 50 defective units. This is not a recommendation: it illustrates what an incomplete financial objective can permit.

Proposition 8.2 (History-specific release gates compose)**.** If a policy releases only at histories k where $\mathbb{P}(D > 50 \mid k) \le \varepsilon$, then

$$\mathbb{P}(\text{release}, D > 50) \le \varepsilon \mathbb{P}(\text{release}). \tag{8.4}$$

The statement holds under the probability model used to evaluate the gate. It does not establish model calibration.

At $\varepsilon = 0.01$, the selected menu option purchases 200 coupons and releases only for $k \le 5$. Expected cost is CA$585,913.86, versus CA$950,000 for the no-test risk-feasible alternative of automatic rework: a conditional difference of CA$364,086.14 per authored shipment. Release probability is 66.02%; joint unsafe-release probability is 0.1472%. This is a model-dependent value of testing, not customer Alpha or a universal inspection optimum.

9 The inspector must itself be measured

The coupon likelihood is only useful if its sensitivity, false-alarm rate and dependence structure are applicable. The extension challenges the frozen nominal policy with mean defects 0.04, lot correlation 0.03, sensitivity 0.75 and a 2% whole-test blindness event. In that event all coupon flags are zero regardless of the latent defect fraction. The policy is not allowed to refit after seeing the challenge.

Its joint unsafe-release probability rises to 1.849%, and the conditional risk among released shipments is about 3.63%. The nominally valid gate is no longer valid under the changed measurement world. A local mathematical proof about a posterior cannot prevent the posterior's likelihood from being wrong.

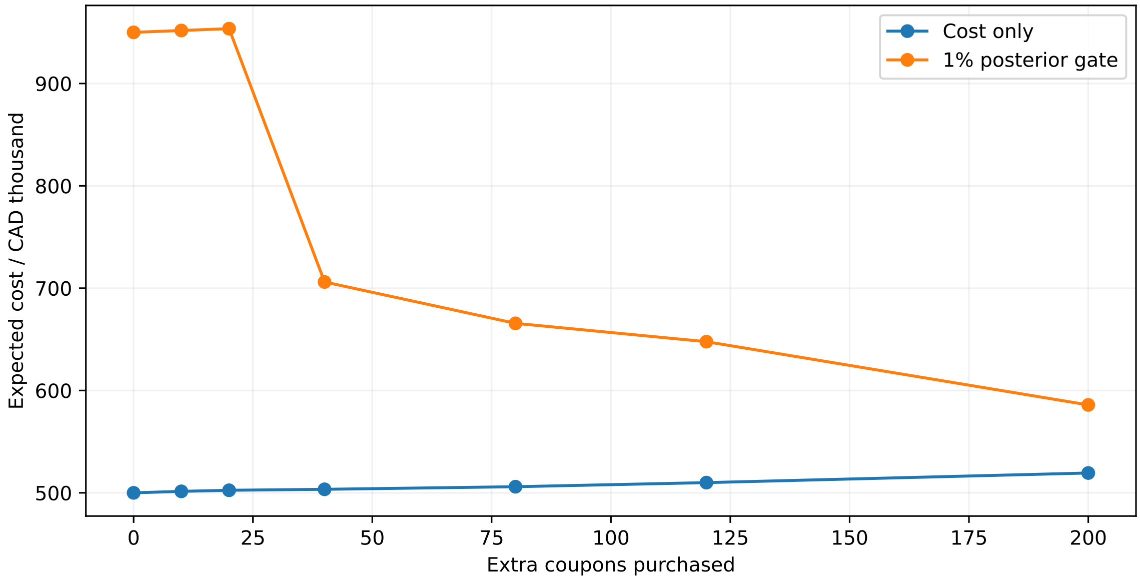

Figure 4. Expected cost across the permitted coupon menu. The gated and ungated policies solve different objectives; the cheaper ungated answer does not meet the declared release-risk condition.

The sensitivity grid contains 27 combinations of mean defect rate, dependence and coupon cost. Fourteen choose no further coupons under the risk gate; some release, others rework. No single sample size is promoted as a best practice across those worlds. The owner should purchase a different measurement channel when it changes the decision or reduces a binding risk, not merely because it produces another report.

Required real-world falsifiers

Reference-material disagreement; blinded duplicate disagreement; detector sensitivity drift; simultaneous failures across instruments; changes in lot formation; missing chain of custody; unexplained label inflation. Each should trigger a scoped loss of reliance, not a cosmetic confidence adjustment.

10 Research value and economically relevant knowledge

For a belief state π and permissible actions a, define $V(\pi) = \max_a \mathbb{E}_\pi U(a, \theta)$. A proposed experiment e with observation Y and complete cost c_e has one-step decision value

$$\mathrm{EVSI}(e; \pi) = \mathbb{E}_Y V(\pi_{e,Y}) - V(\pi) - c_e. \tag{10.1}$$

Without experiment costs the first difference is nonnegative when the owner can ignore information. It can nevertheless be exactly zero for a scientifically informative observation if the same action remains optimal at every posterior. A detector that changes a release decision can be worth more than a more accurate predictor in an economically irrelevant region. The new inspection study is an explicit finite example.

A multi-experiment programme replaces this expression by a finite-horizon Bellman recursion with a stop action, costs, budget, time and risk constraints. That extension is a research formulation here, not a newly executed autonomous experimental laboratory. v10.2's neural disagreement score is a useful proposal mechanism, but the DAEDALON ablation does not show that it maximizes expected value of information.

The useful owner strategy is thus not to maximize model complexity. It is to allocate research effort to the decision boundary, while preserving enough exploratory diversity to detect a missing mechanism. A new method that cannot beat informed Beta may still supply reusable observations or falsifiers, but their value must be measured independently of an unsupported prediction premium.

11 From normalized response to physical process models

The executable acceptance model is a statistical process representation. Its coefficients are not heat-transfer coefficients, activation energies or material constants. A field implementation would need to establish how normalized controls map to qualified instruments and physical states. The following equations define a prospective physical interface, not a physics solver implemented by the current customer adapter.

11.1 Energy and material interfaces

A continuum thermal submodel can start with

$$\rho c_p(\partial_t T + \mathbf{v} \cdot \nabla T) = \nabla \cdot (\kappa \nabla T) + \dot{q}_{\text{electric}} + \sum_j (-\Delta h_j) r_j, \tag{11.1}$$

with declared density ρ, heat capacity c_p, conductivity κ, velocity $\mathbf{v}$, electrical volumetric heating and reaction rates r_j. A radiating boundary requires a consistent heat balance, for example

$$-\mathbf{n} \cdot \kappa \nabla T = h(T - T_a) + \epsilon\sigma(T^4 - T_s^4) - \alpha_s I_s. \tag{11.2}$$

A scalar emissivity and uniform temperature are simplifying assumptions. Spectral absorptivity and emissivity, view factors, orientation, heat transport, surface aging and incident flux affect the applicability of this reduction. The external NASA reference provides thermal-engineering context; it does not validate the invented DAEDALON factory.[9]

The observation interface should preserve dimensions and balances before statistical fitting. A normalized dose might depend on dwell time, power and characteristic material scales; changing the normalization creates a changed model contract. A temperature sensor reporting Celsius into a Kelvin fourth-power law is not a small statistical perturbation. The evidence fabric must retain units, calibration and transformations rather than flattening measurements into undifferentiated numbers.

11.2 A mesoscopic route, not a metaphor

For discrete material states with probabilities x_i and bidirectional rates k_{ij}, a master equation represents unresolved transitions. Its statistical entropy production is

$$\dot{S}_{\text{prod}} = \frac{k_B}{2} \sum_{i,j} (k_{ij} x_i - k_{ji} x_j) \log \frac{k_{ij} x_i}{k_{ji} x_j} \geq 0, \tag{11.3}$$

for positive forward/reverse fluxes, with limits treated appropriately. The inequality follows termwise from $(a - b)\log(a/b) \geq 0$. Interpreting the expression as physical heat/entropy requires thermodynamically compatible rates and specified reservoirs, not merely a fitted graph. This is a route to constrained mechanistic learning inspired by master-equation thermodynamics. It is not a law that an economic institution's wealth must increase.[10]

12 Energy-aware control and thermal invariance

Port-Hamiltonian modeling organizes energy storage, dissipation and interconnection without treating every subsystem as an unconstrained predictor. A finite-dimensional representation is

$$\dot{x} = [J(x) - R(x)]\nabla H(x) + G(x)v, \quad J = -J^\top, \quad R \succeq 0, \quad y = G^\top \nabla H. \tag{12.1}$$

Thus $\dot{H} = -\nabla H^\top R \nabla H + y^\top v$. This supplies an energy-accounting structure, not a guarantee of safe operation or asymptotic stability without further assumptions. Network interconnections can preserve the supply relation when boundary variables and Dirac interconnections are correctly defined.[11, 12]

Proposition 12.1 (A thermal availability function)**.** For $C > 0, T, T_a > 0$ and

$$C\dot{T} = v - k(T - T_a) - \epsilon\sigma A(T^4 - T_a^4), \tag{12.2}$$

with $k, \epsilon, A \geq 0$, define

$$\mathcal{B}(T) = C[T - T_a - T_a \log(T/T_a)] \geq 0. \tag{12.3}$$

Then

$$\dot{\mathcal{B}} = (1 - T_a/T)v - \frac{k(T - T_a)^2}{T} - \frac{\epsilon\sigma A(T - T_a)(T^4 - T_a^4)}{T} \leq (1 - T_a/T)v. \tag{12.4}$$

The two dissipative terms are nonpositive on either side of ambient temperature. This physically interpretable check complements a learned response model. It cannot compensate for an unmodeled heat source, time-varying heat capacity or invalid radiative approximation.

Proposition 12.2 (A robust upper-temperature gate)**.** Suppose $C\dot{T} = v + d - \ell(T)$, $d \leq \bar{d}$, and the closed-loop solution is well posed. If an enforced controller satisfies

$$v + \bar{d} - \ell(T) \leq C\gamma(T_{\max} - T), \qquad \gamma > 0, \tag{12.5}$$

and $T(0) \leq T_{\max}$, then $T(t) \leq T_{\max}$ for as long as the assumptions remain valid.

Let $h = T_{\max} - T$; the inequality gives $\dot{h} \geq -\gamma h$ and comparison yields $h(t) \geq e^{-\gamma t}h(0)$. The controller must have a feasible action satisfying the inequality; otherwise it needs a separately validated shutdown or safe-mode path. State-estimation error, sample delay and actuator saturation require additional margins. This is the role of a control barrier, not a proof obtained by describing a desired boundary in prose.[13]

No such apparatus controller is activated by this publication. These equations identify interfaces and proof obligations for a future mission adapter. The separation from the current finite synthetic planner is essential to a useful assurance case.

13 Strategic quality: make truthful production worth doing

A manufacturing institution is also a game among the owner, process operator, model supplier, instrument provider and verifier. A profitable deviation may be to skip an expensive step, relabel defects, choose an easier test or hide a failed lot. Physical consistency checks and economic incentives address different failure modes. Strategic inspection has a long mathematical tradition; the present contract is a deliberately small industrial specialization.[14]

Proposition 13.1 (Imperfect-inspection deterrence)**.** A risk-neutral operator gains $g > 0$ by taking low effort. Conditional on an audit, a sanctionable signal occurs with probabilities d_L under low effort and d_H under high effort. Audits occur with probability a, enforcement succeeds with probability κ, and an enforceable loss is B. With no other payoff differences, high effort is a best response exactly when

$$a\kappa(d_L - d_H)B \geq g. \tag{13.1}$$

False positives matter because d_H is subtracted, not ignored. Limited liability bounds B; if the inequality cannot hold even at feasible maximum audit and sanction levels, the contract cannot induce the intended effort. The response is to redesign the process or contract, not to increase a nominal bond beyond what can be collected. Audit costs, wrongful-sanction costs, appeals and the operator's participation constraint must also be budgeted.

Common control of the operator and final verifier defeats the intended mechanism. An instrument can be technically separate while its incentives, evidence access and funding remain correlated. Real independence is established by custody, governance, conflicts and test control; multiple signatures alone do not supply it.

14 Truthful forecasts are not enough: pay for useful evidence

A supplier asked to report the probability of a defect can be scored against an externally resolved outcome. For $Y \in \{0,1\}$ and report r, the quadratic score $S(r,Y) = -(r-Y)^2$ is strictly proper: at belief p, expected score is $-(r-p)^2 - p(1-p)$ and is uniquely maximized by $r = p$. Proper scoring supplies an incentive for honest reporting under the stated conditions; it does not itself pay for acquiring better information.[15]

Proposition 14.1 (The evidence-acquisition incentive)**.** Let prior $p = \mathbb{E}Y$, and let a paid signal Z produce posterior $p_Z = \mathbb{E}[Y \mid Z]$. If the supplier truthfully reports and is paid wS, the expected score gain from the signal is

$$w\{p(1-p) - \mathbb{E}[p_Z(1-p_Z)]\} = w\,\mathrm{Var}(p_Z). \tag{14.1}$$

Acquisition is weakly preferred to no acquisition when $w\,\mathrm{Var}(p_Z) \geq c_Z$, absent other payoff effects.

The identity is the law of total variance. It exposes a design tension: a supplier may be completely honest yet rationally choose not to perform costly useful measurements. Conversely, a large score multiplier can reward information that is irrelevant to the customer's decision. A practical contract combines truthful outcome scoring with a preconstituted experimental brief, paid effort where appropriate, a proper decision objective and independent outcome custody.

The theorem assumes the report does not change the scored event. If a forecast triggers a maintenance action that prevents the predicted failure, naïvely scoring the observed nonfailure can penalize a good forecast. Action-conditioned targets, protected evaluation cases or explicitly causal evaluation are then required. This is not resolved by adding another language-model critic.

15 Coordinating scarce metrology without inventing capacity

Inspection creates its own bottleneck. Let each of finitely many jobs choose a subset of laboratory resources. If resource ℓ has load n_ℓ and an identical per-user cost $c_\ell(n_\ell)$, a job pays the sum of those costs plus its private route cost $b_i(a_i)$. The unit-load assumption matters: weighted jobs, player-specific congestion and coupled capacity constraints are not automatically covered.

Proposition 15.1 (A finite coordination potential)**.** Under those assumptions,

$$\Psi(a) = -\sum_i b_i(a_i) - \sum_\ell \sum_{k=1}^{n_\ell(a)} c_\ell(k) \tag{15.1}$$

is an exact utility potential. Any strict unilateral cost improvement raises Ψ, so finite better-response dynamics terminate at a pure Nash equilibrium.

Only the terms on resources entered or left change, giving exactly the deviator's utility change. The equilibrium need not minimize total social cost. Marginal congestion tolls, joint scheduling or revised contracts may improve the allocation; no computed equilibrium creates an additional laboratory slot. Potential-game theory supplies the formal structure, not an empirical law of industrial cooperation.[16]

Proposition 15.2 (The exact finite Gibbs connection)**.** If one player at a time is selected uniformly and updates with logit probability proportional to $\exp(u_i/\tau)$, $\tau > 0$, in a finite exact potential game with all actions available, the stationary distribution is

$$\pi_\tau(a) = \frac{e^{\Psi(a)/\tau}}{Z_\tau}, \qquad \max_\nu\{\mathbb{E}_\nu\Psi + \tau H(\nu)\} = \tau \log Z_\tau. \tag{15.2}$$

Detailed balance follows from the equality of payoff and potential changes. The variational gap is $\tau D_{\mathrm{KL}}(\nu\|\pi_\tau)$. This is a precise mathematical connection with finite statistical mechanics, not an assertion that an institution has a physical temperature. Slow mixing, synchronous updates and nonpotential incentives can invalidate the useful equilibrium interpretation.[17, 18]

Correlated recommendations can also coordinate different laboratories, provided the recommendation distribution satisfies incentive constraints and respects actual slots. Correlated actions are sometimes useful; correlated evidence is not independent evidence. If a common-mode defect occurs with probability q, a committee's failure can have a floor at q even when its remaining errors are independent. Repeating the same detector cannot remove that floor.

16 A comparison budget across successive candidates

A succession programme can repeatedly fit candidates and inspect results until a favourable one appears. A conventional interval for one fixed comparison does not automatically cover that process. Let $D_{gjt} \in [-B, B]$ be a paired utility difference for frozen generation g, comparator $j \in \{1, \dots, m\}$ and observation t. Conditional on the prior history, suppose each frozen pair has a stationary mean μ_{gj} and satisfies the bounded concentration assumptions.

Proposition 16.1 (A conservative successor-wide certificate)**.** With $\alpha \in (0, 1)$, simultaneously for all g, j, t, with probability at least $1 - \alpha$,

$$\mu_{gj} \geq \overline{D}_{gjt} - B\sqrt{\frac{2\log[mg(g+1)t(t+1)/\alpha]}{t}}. \tag{16.1}$$

Apply one-sided Hoeffding with error allocation $\alpha/[mg(g+1)t(t+1)]$, then union-bound over the countable indices. This deliberately conservative construction illustrates the need to budget repeated looks; sharper confidence sequences exist.[19, 20] Conditional independence from fresh cases and a fixed release are indispensable. Unbounded losses, adaptive reuse of protected cases or an evolving mean require different methods.

For admission, the lower bound must exceed the declared margin against every required comparator after incremental costs, human burden and risk/impairment reserves, while all hard gates pass. This paper does not retroactively apply that certificate to the eight synthetic cohorts: their bootstrap is retained as a conditional analysis with a different design. Statistical confidence, customer acceptance and authority remain different objects.

17 The constitutional invariant of succession

A deployment claim belongs to a tuple: mission, exact release, evidence commitment, protocol, comparator set, verifier, operating environment and validity window. A changed tuple requires an explicit transfer decision. This follows the source architecture's proof-currency doctrine and its distinction between a useful candidate and an admitted SUCCESSOR Ω.[2]

No inherited permission

G3 may inherit observations, negative results, interface contracts and its parent's identity link. It may not inherit the parent's proof receipt, safety conclusion, economic superiority or right to actuate equipment.

An implementation enforces this by binding each authorized request to an exact release and checking a current, nonrevoked receipt for that release. If the admission service has no qualifying receipt for the new identifier, the request fails. This is an engineering invariant under trusted identity, policy and execution enforcement, not a theorem that software cannot be compromised. The current manufacturing extension invokes no such physical authorization service.

The owner programme must separately qualify source rights, calibrated instruments, protected custody, supplier substitution, threat controls, human accountability, recovery and monitoring. The source's A0–A4 ladder describes increasing scoped authority; the artifacts here establish research computation only. NIST's AI risk framework is relevant governance context, not a certification of this product.[21]

Replace what wears out.
Then expand what is possible.

Capacity is an age-structured stock, not a perpetual promise.

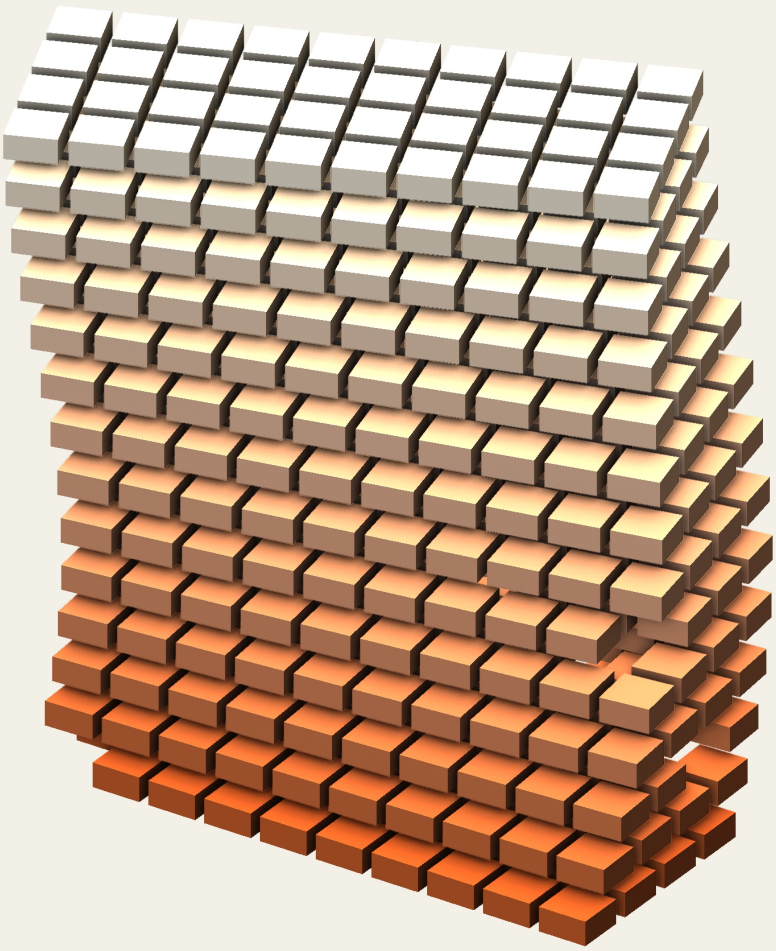

Cash pays for commissioning. Material feeds replacement.
No acquired knowledge cancels those obligations.

Conceptual cohort sculpture, not a plotted result or a designed factory.

18 Accepted production, paid construction and delayed commissioning

The retained capital continuation is a separate business scenario, not another training trial. It uses mean own-production quantities and expense ratios from the second mission as stationary coefficients. That deliberately discards plant-level heterogeneity and lifetime production uncertainty; it must not be mistaken for a stochastic robust financing policy.

Both routes start with eight factories, two metrology laboratories and CA$1.6 billion cash. A new factory requires CA$1.1 billion, 80,000 accepted internal panels and three years. A laboratory costs CA$70 million and takes two years; each serves four factories. Both asset classes retire at 20 years. The panels are only one illustrative bill-of-materials component; all remaining machinery and integration are purchased in the cash cost. No self-replicating universal factory is implemented.

$$C_{t+1} = C_t + R_t - O_t - I_t - D_t - S_t, \tag{18.1}$$

$$G_t^{\text{own}} = G_t^{\text{external}} + G_t^{\text{internal}} + G_t^{\text{held}} + G_t^{\text{unused}}, \tag{18.2}$$

$$K_t = \sum_v x_v \mathbf{1}\{v + \tau \le t < v + \tau + L\}. \tag{18.3}$$

Here C is cash; R recognized cash receipts; O operating and reserve outflows; I paid investment; D owner distributions; S science allocations; and x_v the asset cohort ordered at time v. The good-unit partition is a per-period disposition without double counting; carried inventories need their own opening and closing account.

Factory starts are capped at three per year, laboratory starts at two. Electricity is capped at 450 GWh/year and feedstock at 1.2 million tonnes/year. Demand grows from 3.5 to 10.5 million panels/year and price declines 0.5% annually. Positive operating cash retains 25% as an authored tax/working-capital allowance; of the remainder, 15% is distributed and 10% funds science. These are scenario conventions, not tax advice.

Years 9 and 17 impose lower availability and CA$18 million requalification cost. A CA$600 million cash floor, idling and standalone laboratory replacement are supported. The investment selector enumerates local choices over a 12-year underwriting window at 8% discount; it does not prove a global 24-year optimum. Paying requalification cost in the model does not itself establish that a real process became safe.

19 The 24-year account: more factories, not automatic dominance

Scenario endpoint	Fixed	Candidate	Informed Beta
Operating factories	18	22	22
Metrology laboratories	6	7	7
External panels sold, million	162.454	168.483	168.483
Construction paid, billion CAD	22.490	26.960	26.960
Owner distributions, billion CAD	15.832	16.183	16.202
Science funding, billion CAD	10.555	10.789	10.801
Closing cash, billion CAD	58.270	55.548	55.649

Relative to fixed operation, the candidate route has four more operating factories, CA$350.793 million more cumulative owner distributions and CA$233.862 million more science funding. It also has CA$2.722 billion less closing cash after CA$4.470 billion more paid construction. These undiscounted quantities do not establish a larger enterprise valuation without a continuation or liquidation convention.

Informed Beta obtains the same physical path under lower programme overhead. Its extra cash, distributions and science allocations total CA$132 million relative to the candidate, corresponding to the modeled post-reserve treatment of the candidate's CA$174 million total incremental programme costs. Most growth comes from common margin and demand assumptions, not from a proprietary model advantage.

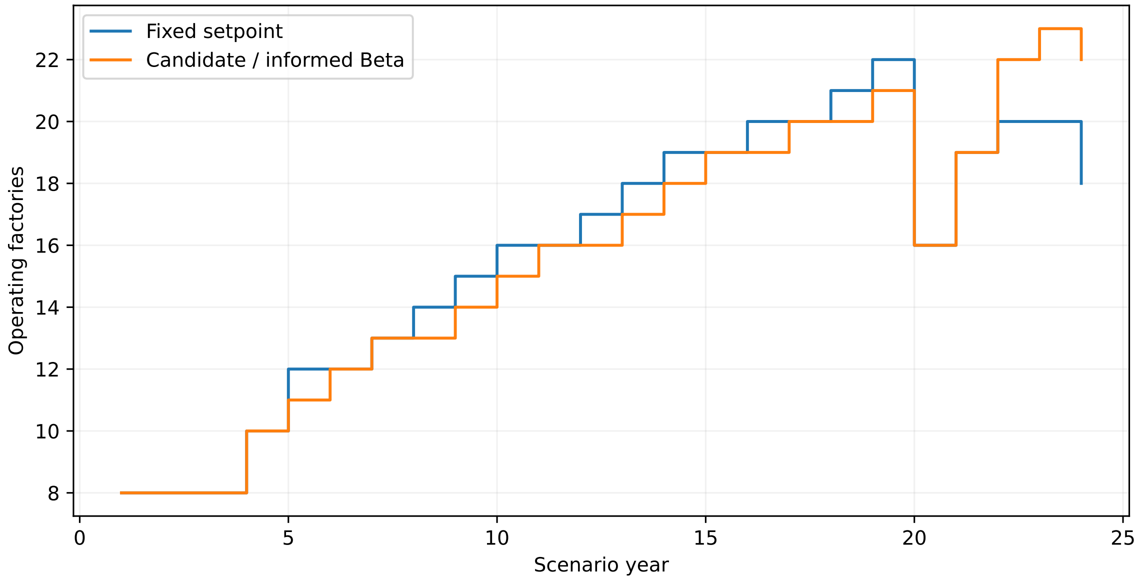

Figure 5. Operating assets in the conditional capital account. Informed Beta and candidate have the same physical path; offsets must not be invented to separate overlapping curves.

The retained sensitivity analysis contains 81 paths; those paths complete under their selected ranges. A separate lower-price case stops at the cash floor. The earlier exploratory transfer of fulfillment settings into an own-sales business is also preserved. Those failed or revised exploratory choices are not converted into confirmatory evidence by the final presentation.

20 A network that can reproduce productive capability

A scalar factory count omits complementary capabilities. Fabrication requires instruments, energy, feedstock, assembly, transport and repair. A prospective age-structured model uses a nonnegative kernel $K_{ij}(a)$: the expected rate of newly qualified type-i capability generated by one type-j cohort of age a, after quality, survival and commissioning effects. Then births obey

$$b_i(t) = \sum_j \int_0^\infty K_{ij}(a) b_j(t-a)\,\mathrm{d}a. \tag{20.1}$$

This is a proposed reproduction model, not a measured industrial law. It becomes useful only after coefficients are identified, units fixed, and cash/material/energy constraints enforced outside the unconstrained kernel. A predictive model does not fabricate its own instruments merely because those instruments appear as nodes.

Proposition 20.1 (A conditional productive-reproduction threshold). Assume an integrable nonnegative kernel with irreducible Laplace transform, reproduction only at positive ages, and strict decay of its spectral radius with g. A balanced positive-growth solution $b(t) = e^{gt}v$, $v > 0$, exists with $g > 0$ if and only if

$$\rho\left(\int_0^\infty K(a)\,\mathrm{d}a\right) > 1; \quad \rho\left(\int_0^\infty e^{-ga} K(a)\,\mathrm{d}a\right) = 1 \tag{20.2}$$

determines the growth rate. This is an unconstrained local reproduction condition, not a claim of indefinite feasible expansion.

Substitution gives a Perron eigenvalue equation; continuity and strict monotonicity give the crossing. For the scalar kernel $K(a) = \beta e^{-\delta a}\mathbf{1}\{a \geq \tau\}$,

$$1 = \frac{\beta e^{-(g+\delta)\tau}}{g+\delta}, \qquad g > 0 \iff \beta e^{-\delta\tau} > \delta. \tag{20.3}$$

A commissioning delay reduces reproductive value even before cash discounting is applied. Metrology that improves accepted quality can raise an element of K; a common calibration failure can lower several elements simultaneously. A credible growth programme must therefore measure the weak links, not multiply independent success probabilities it has never established.

21 Replacement before expansion

An installed estate must be maintained. Let $B(t)$ be installed mass per year, let $S(a)$ be the surviving fraction at age a, and let the total surviving mass be

$$M(t) = \int_0^\infty B(t-a)S(a)\,\mathrm{d}a. \tag{21.1}$$

On a balanced exponential path $B(t) = B_0 e^{gt}$, throughput per unit stock is $B/M = [\int e^{-ga}S(a)\,\mathrm{d}a]^{-1}$. It is a stock-flow relationship, not an economic return.

Proposition 21.1 (Deterministic-life replacement and growth). If every installed unit survives exactly L years and balanced growth is $g \geq 0$,

$$\frac{B}{M} = \begin{cases} L^{-1}, & g = 0, \\ g/(1-e^{-gL}), & g > 0, \end{cases} \qquad R = B - gM, \qquad V = B - rR = gM + (1-r)R, \tag{21.2}$$

where R is retired mass and $r \in [0,1]$ the qualified recovery fraction, so V is required virgin feedstock.

Even perfect recovery cannot supply the net new stock gM without new material. Recycling may reduce virgin input but does not eliminate gross fabrication, handling, inspection, repair energy or transportation. A 50-year stock growing at 2% requires 1.582 times its zero-growth replacement throughput; at 5%, the factor is 2.724. These conditional calculations are included in the new laboratory.

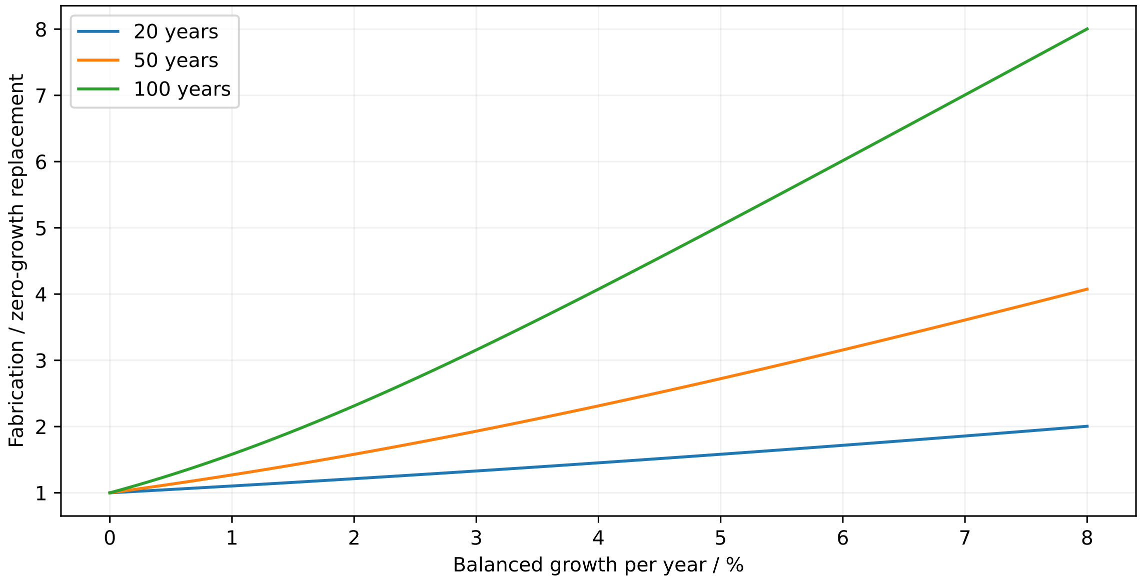

Figure 6. Gross fabrication needed for a balanced growing stock, divided by zero-growth replacement. Lifetime and growth are assumptions, not measured projections.

22 The stellar envelope: power, heat, matter and time

The long horizon is a scientific and productive network that can support progressively larger missions. It is not a forecast that the 24-year financial model reaches stellar scale. The physical laboratory separates useful power P_u from captured incident power $P_c = P_u/\eta$. At nominal one-AU irradiance $I_\odot = 1361\ \mathrm{W/m^2}$ and nominal solar luminosity $L_\odot = 3.828 \times 10^{26}$ W,[22]

$$A_c = \frac{P_c}{I_\odot}, \quad A_r = \frac{P_c}{\epsilon\sigma(T^4 - T_s^4)}, \quad M = \mu(A_c + A_r), \quad P_c \le L_\odot. \tag{22.1}$$

The radiator expression assumes all captured power eventually thermalizes inside the chosen boundary. Exporting useful energy changes that boundary; claiming only conversion waste while retaining all useful power inside would undercount eventual heat. Projected collecting area is not the complete area of structures or a deployable geometry.

Illustrative input or consequence	Value
$P_u, \eta, T, \epsilon, T_s$	10^{26} W, 0.30, 600 K, 0.90, 3 K
Captured power / nominal solar luminosity	87.08%
Projected collection area	$2.449 \times 10^{23}\ \mathrm{m^2}$
Effective emitting area	$5.040 \times 10^{22}\ \mathrm{m^2}$
Film-only mass at assumed $\mu = 0.1\ \mathrm{kg/m^2}$	2.953×10^{22} kg
50-year, zero-growth replacement	5.906×10^{20} kg/yr
50-year stock at 2% growth: gross fabrication	9.344×10^{20} kg/yr
Same path, 90% recovery: virgin material	6.250×10^{20} kg/yr

These quantities omit structures, mutual shadowing, optical constraints, transport, reprocessing energy, thermal plumbing, redundancy, radiation damage and correlated failures. At $\eta = 0.20$, the same useful-power target exceeds a single nominal star's output and the laboratory flags it. Finite light speed also matters: one AU takes about 499 seconds one way, before computation and communication overhead. Fast safety functions must be local or preauthorized; a distant executive cannot close an impossible control loop.

Radiation pressure adds a momentum budget. At normal incidence, absorption produces approximately I/c pressure and perfect reflection $2I/c$; the resulting force depends on orientation and optical properties. Light sheets are not massless, force-free structures. NASA's ARMADAS programme provides external evidence that modular autonomous assembly is a genuine research direction, but it is not integrated with DAEDALON and does not establish this scale.[23]

23 The highest-value owner programme

The most attractive use is not an unconditional purchase of model superiority. It is a customer-owned mission institution able to obtain the best suitable predictor, preserve hard-won evidence, qualify physical outcomes and allocate investment to the next binding constraint. The owner can keep the institutional knowledge while selecting informed Beta whenever that choice is cheaper and equally capable.

Gate	Owner decision	Evidence that would justify proceeding
D0	Select one recurring costly decision.	Observable output, lawful evidence, meaningful comparator and finite exposure.
D1	Purchase better understanding.	Instrument qualification, baseline reproduction, registered objective and falsifiers.
D2	Freeze a candidate for independent challenge.	Complete manifest, protected cases, costly failure modes and current economics.
D3	Grant the least authority that earns value.	Independent result, operational assurance, named responsibility and recovery drill.
D4	Renew or replace the institution's composition.	Current metrology, contemporary Beta, drift tests, retirement and cash obligations.

A defensible premium could come from a genuinely different physical process, access to lawful proprietary evidence, better instrument calibration, faster justified impairment detection, lower whole-system cost or demonstrated recovery quality. None of those benefits follows from the present model name. A CA$7 million annual overhead needs more than CA$7 million of incremental conservatively supported benefit before a positive margin and nonfinancial hard gates are considered.

The proposed route to an admitted Specialist ASI is mission-specific: precise equipment family, process range, instruments, objective, counterparties, current baseline and validity period. The delivered candidates have not earned that admission. If future work does, the customer acquires a bounded operating institution with evidence and renewal obligations, not an unrestricted scientific oracle. The architecture's claims-to-proof covenant is the governing standard.[2]

24 Limitations and the productive thesis

The synthetic mechanism is grammar-compatible and known to the authors. Coupon means have unusually simple bounded noise. Conditional binomial lot noise omits real defect correlations; the new beta-binomial and metrology examples expose this limitation but do not repair or requalify the original industrial study. The finite comparator set is not a market survey. Same-author custody, eight refits and local deterministic replay do not create independent proof.

Capital margins, taxes/working-capital reserves, customer demand, material prices, asset lives and bottlenecks are assumptions. Stationary mean production coefficients cannot establish financing resilience under correlated failures. The rolling investment routine is not globally optimal, nor is the 24-year endpoint an enterprise valuation. The stellar sheet-mass calculation omits most of a real system. No scientific result supplies a timetable for constructing it.

Those limitations do not make the institution useless. They identify what is worth learning next. The executed study demonstrates that explicit world programs can be formed, contradicted and revised, and that a changed objective changes the productive decision. The new calculations show that metrology, dependence and replacement must be first-class parts of the mission world. Strategic contracts must reward real effort and honest evidence, while preserving the owner's ability to reject the current candidate.

The Capacity Foundry

Knowledge becomes a productive asset only when it changes a defensible decision. Productive assets become an enduring estate only when accepted cash, qualified replacements and accountable succession keep them useful. The grand ambition is not endless confidence. It is a continuously renewable capacity to discover, manufacture and preserve what reality permits.

A Notation, units and noninterchangeable quantities

Symbol	Unit	Meaning
u, z	1	Normalized intensity and contamination; not apparatus settings.
p, G, n	1; units	Mean acceptance, accepted count, attempted count.
q, D	1; units	Latent defect fraction and shipment defect count in the new inspection study.
E, M	kWh; kg	Electrical energy and material mass.
C_t, R_t, I_t	CAD	Cash, received cash, paid investment; not valuation.
P_u, P_c	W	Useful and captured power, not heat-rejection capacity.
A_c, A_r	m^2	Projected collector and effective emitting areas.
$B(t), R(t), V(t)$	kg/yr	Gross installation, retirement, virgin input in the renewal model.
g, L	yr^{-1}; yr	Balanced growth rate and deterministic life.
τ	Context dependent	Commissioning delay in renewal; a payoff-scale parameter in logit. Never a physical temperature there.
α, ε	1	Statistical error budget and posterior release-risk threshold.

Annual financial contrasts use the stipulated 13,312 line-weeks. The capital account uses 24 undiscounted scenario years. The stellar envelope is a physical arithmetic exercise. These domains cannot be combined by summing their attractive headline numbers.

B Complete formation and comparison protocol

Stage	Cases per cohort	Use
G1 training / selection	32 / 20	Narrow domain
Permitted neural-guided probe pool	160	One selected
G2 new training + selected probe	128 + 1	Development
G2 selection / holdout / stress	64 / 128 / 64	Separate IDs
Decision-margin selection	48	Common cases
Nominal / shifted / supply-shock decisions	128 / 64 / 32	Paired draws
G3 training / selection	192 / 64	Shifted domain
G3 holdout / stress	96 / 64	Fresh local IDs
Own-sales mission	128	Frozen models

Primary cohort seeds are 1100–1107. The source preserves a two-cohort pilot and an exploratory objective-transfer history. Native search starts with maximum polynomial degree one, at most 80 basis expressions, four terms, beam width 12, at most 3,000 fits and a 60-second budget per declared search. Adaptive revision permits degree two. Generic Beta fits degrees 1, 2, 3 and 5 under two regularizations, two ExtraTrees variants and two histogram-boosting variants. The informed control is included in every strongest-baseline interpretation.

The inherited protocol hash is retained in the original `PROTOCOL.json`. Reading the hash proves neither author identity nor prospective registration. It identifies the exact local protocol bytes. The publication does not change that protocol, add undisclosed winning trials, or reclassify exploratory adjustments as confirmatory results.

C Cohort-level economic ledger

Cohort	Fixed contrast	Generic edge	After costs	Shift gain/week
0	1172.630	6.661	-0.339	0.012440
1	1256.015	-0.772	-7.772	0.010309
2	1061.325	-8.001	-15.001	0.015304
3	1323.995	10.362	3.362	0.006824
4	1174.926	2.995	-4.005	0.010568
5	1177.736	1.235	-5.765	0.021683
6	1138.754	4.959	-2.041	0.017176
7	1033.957	3.950	-3.050	0.021917

All entries are authored annual CAD millions at the same 32-site scale. The informed-Beta difference before overhead is zero in every nominal cohort. The negative mean after costs is preserved, even though the candidate's generic-Beta pre-cost edge is positive in several cohorts.

D Proofs of capped-output and regret bounds

Proof of Proposition 5.1. For integer $G \geq 0$, $\min(G, d) = \sum_{k=1}^{d} \mathbf{1}\{G \geq k\}$. Taking expectations gives the tail sum. Splitting at $G = d$ and using the binomial coefficient identity gives the CDF formula. More transparently, $k\mathbb{P}(G = k) = np\mathbb{P}(G_{n-1} = k - 1)$. For any function f on $\{0, \ldots, n\}$, differentiating the Bernstein polynomial gives

$$\frac{\mathrm{d}}{\mathrm{d}p}\mathbb{E}f(G_n) = n\mathbb{E}[f(G_{n-1} + 1) - f(G_{n-1})]. \tag{D.1}$$

For $f(k) = \min(k, d)$ the increment is one exactly when $k \leq d - 1$. For $d = 0$ the expectation is zero; for $d \geq n$ it is np. □

Proof of Proposition 5.2. Couple the two binomials using common uniforms U_i: $G_p = \sum_i \mathbf{1}\{U_i \leq p\}$. Then $\mathbb{E}|G_p - G_{\hat{p}}| = n|p - \hat{p}|$. The payoff difference at each action is bounded by $Ln_{\max}e$. Add and subtract approximate utility at the true and approximate maximizing actions; the approximate-optimality term is nonpositive, leaving two error terms. □

The coupling proof requires the specified conditional distribution. It does not show that matching marginal acceptance means suffices when batch dependence changes. That is exactly the gap exposed by the beta-binomial calculation.

E Inspection posterior, quadrature and dependence

Proof of Proposition 8.1. Conditional variance gives $\mathrm{Var}(G) = \mathbb{E}[nP(1-P)] + \mathrm{Var}(nP)$. Since $\mathrm{Var}(P) = p(1-p)\rho$, substitution yields $np(1-p)[1+(n-1)\rho]$. Dividing by n^2 and matching the variance of an independent sample mean yields n_{eff}. □

For the coupon study, the beta prior has shape parameters $a = 1.98$, $b = 97.02$. The instrument is imperfect, so the posterior is not generally beta-conjugate:

$$\pi(q \mid K = k) \propto q^{a-1}(1-q)^{b-1} h(q)^k [1-h(q)]^{m-k}. \tag{E.1}$$

The release gate integrates the binomial upper tail of a separate shipment conditional on the same q. The code uses 640-node Gauss–Jacobi quadrature; for the declared 1,000-unit shipment and at most 200 coupons, the relevant polynomial factors have degree at most 1,200. The quadrature is degree-exact in ideal arithmetic; the implementation is floating point and is compared with 768 nodes and direct beta-binomial no-test tails. It is not interval arithmetic or a proof of every numerical backend.

Proof of Proposition 8.2. Sum $\mathbb{P}(K = k)\mathbb{P}(D > 50 \mid k)$ over released histories only. Each conditional risk is at most ε, so the sum is at most $\varepsilon \sum_{k:\text{release}} \mathbb{P}(K = k)$. The result is a joint-risk statement and, when release probability is positive, also bounds risk conditional on release. □

The impaired detector mixture is $(1-b)\,\mathrm{Binomial}(m, h(q)) + b\delta_0$ with $b = 0.02$. Multiplying individual detector success probabilities would miss this common blindness event.

F Thermal and strategic derivations

Proof of Proposition 12.1. Write $x = T/T_a > 0$. Then $\mathcal{B} = CT_a(x - 1 - \log x) \geq 0$ by convexity. Differentiate: $\dot{\mathcal{B}} = C(1 - T_a/T)\dot{T}$. Substituting the heat balance gives the displayed identity. Both $(T - T_a)^2$ and $(T - T_a)(T^4 - T_a^4)$ are nonnegative. □

Proof of Proposition 13.1. Low effort minus high effort changes expected payoff by $g - a\kappa(d_L - d_H)B$. Its nonpositivity is equivalent to high effort being a best response. A strict inequality gives a strict preference. This comparison alone does not ensure participation, enforceability, solvency or correct signal probabilities. □

Proof of Proposition 14.1. The optimized expected quadratic score without information is $-\mathrm{Var}(Y)$. With signal Z, it is $-\mathbb{E}\,\mathrm{Var}(Y \mid Z)$. Their difference is $\mathrm{Var}(\mathbb{E}[Y \mid Z]) = \mathrm{Var}(p_Z)$ by total variance. Multiply by payment scale w and compare to the acquisition cost. □

G Finite coordination and repeated proof

Proof of Proposition 15.2. Two profiles differing only in player i have a common conditional logit denominator. Their forward/reverse transition ratio is $\exp([u_i(a') - u_i(a)]/\tau) = \exp([\Psi(a') - \Psi(a)]/\tau)$, which is the ratio of their proposed stationary masses. Detailed balance and irreducibility give stationarity and uniqueness. For any distribution ν,

$$\tau \log Z_\tau - \mathbb{E}_\nu \Psi - \tau H(\nu) = \tau \sum_a \nu(a) \log \frac{\nu(a)}{\pi_\tau(a)} \geq 0. \tag{G.1}$$

□

The numerical check uses a four-profile potential $[0, 3, 3, 2]$ and verifies the transition matrix and variational gap directly. It demonstrates the finite identity, not convergence of real industrial actors to that distribution.

Proof of Proposition 16.1. For differences in $[-B, B]$, one-sided Hoeffding yields $\mathbb{P}(\bar{D} - \mu > r) \le \exp[-tr^2/(2B^2)]$. Choose the displayed r so that the right-hand side is $\alpha/[mg(g+1)t(t+1)]$. The telescoping identity $\sum_{k\ge1} 1/[k(k+1)] = 1$ and the union bound give total error at most α. Apply the argument conditionally at each frozen generation and integrate over prior history. □

A positive bound does not waive nonstatistical gates. A negative bound is not a theorem that no better method can exist; it is insufficient evidence for this exact comparative claim under this test.

H Renewal, conservation and growth limits

Proof of Proposition 20.1. Insert $b(t) = e^{gt}v$ in the renewal equation and cancel e^{gt}. The resulting matrix equation is $\widehat{K}(g)v = v$. Perron–Frobenius gives a positive eigenvector at spectral radius one under irreducibility. The stated monotonicity and decay to zero give a unique positive crossing precisely when $\rho(\widehat{K}(0)) > 1$. This establishes the balanced mode; it does not prove convergence from every age distribution or survival under external constraints. □

Proof of Proposition 21.1. With deterministic survival, $M(t) = \int_0^L B(t)e^{-ga}\,\mathrm{d}a$. The integral is $B(t)L$ for $g = 0$, and $B(t)(1-e^{-gL})/g$ otherwise. Retirement is $R(t) = B(t-L) = B(t)e^{-gL}$, so $B - R = gM$. Qualified recovery rR displaces that amount of virgin material, giving $V = B - rR$. □

For any feasible expansion period, commissioned capacity is bounded by paid investment, accepted internal components, qualified metrology, material, energy, transport and commissioning throughput. Improving one bottleneck cannot be credited as if all others vanished. A factory incapable of replacing its critical instrument is not an indefinitely reproducing productive unit.

I Reproducibility and exact source map

The package includes the original simulation under `simulation/DAEDALON_OMEGA_1_0/`. Its internal manifest, trained artifacts, exploratory records, primary cases, market mission and capital account remain intact. New mathematical code is under `research/`; new paper source is under `latex/`. The new paper is not a runtime version increment.

```
python3 VERIFY_PACKAGE.py
python3 -m pytest research/test_frontier.py -q
python3 research/qualification_frontier.py --out var/new_frontier
cd simulation/DAEDALON_OMEGA_1_0
python3 REPRODUCE_ALL.py --out var/new_replay
python3 -m pytest tests -q
```

Use new output directories. The authoring environment used Python 3.13.5, NumPy 2.3.5, SciPy 1.17.0, scikit-learn 1.8.0 and pytest 9.0.2. Dependency versions are reproducibility metadata, not a vulnerability audit. Dependencies and fonts are not bundled. The reproduction scripts never connect to equipment or send model requests to an external service. Installing dependencies is a separate network operation.

The new mathematical suite checks binomial identities, derivative bounds, correlated variance, posterior gates, quadrature convergence, changed-channel failure, exergy dissipation, incentive inequalities, finite

Gibbs stationarity, renewal balances and stellar capture. The delivered machine-readable validation distinguishes newly executed tests from inherited records. Hashes detect changed bytes relative to a manifest; they are not authenticated publisher signatures.

J The mission qualification packet

A future real engagement should return a decision packet with the following independently checkable objects. Missing entries mean the corresponding claim remains unearned.

Object	Minimum accountable content
Mission and scope	Equipment, objective, permitted actions, exclusions, units, budgets and signed owner.
Measurement qualification	Calibration, sensitivity, specificity, lot dependence, reference custody and uncertainty.
Complete candidate	Model and code hashes, dependencies, limits, frozen prompts where applicable and training record.
Comparative proof	Current procurable Beta, protected cases, preregistered protocol, all costs, tail failures and independent result.
Operating assurance	Threat model, least privilege, physical interlocks, expiry, monitoring, suspend/rollback and recovery drills.
Economic acceptance	Verified delivery, counterparties, actual invoices, cash timing, liabilities and no double counting.
Successor decision	New generation identity, changed claims, fresh required proof and separately granted scope.

No packet in this publication grants real operational authority. Its immediate use is to make the next owner, engineer and verifier decision more precise.

References

[1] MONTREAL.AI. *DAEDALON Ω: The Capacity Foundry, Simulation 1.0.* Supplied research archive. Eight synthetic cohorts; source, observations, trained artifacts, comparison and capital records. Retained in simulation/DAEDALON_OMEGA_1_0. 2026.

[2] V. Boucher. *Neural-Symbolic SUCCESSOR Ω: Customer-Owned Mission Intelligence.* MONTREAL.AI corporate white paper; institutional specification, not a deployment certificate. Supplied source, 59 pages. Aug. 2026.

[3] MONTREAL.AI. *Neural-Symbolic SUCCESSOR Ω v10.2, Discovery Milestone.* Supplied first-party runtime and validation. The retained simulation uses 55 unchanged first-party payload files; no new production admission. 2026.

[4] S. L. Brunton, J. L. Proctor, and J. N. Kutz. « Discovering governing equations from data by sparse identification of nonlinear dynamical systems ». In: *Proceedings of the National Academy of Sciences* 113.15 (2016), pp. 3932–3937. DOI: 10.1073/pnas.1517384113.

[5] S.-M. Udrescu and M. Tegmark. « AI Feynman: A physics-inspired method for symbolic regression ». In: *Science Advances* 6.16 (2020), eaay2631. DOI: 10.1126/sciadv.aay2631. arXiv: 1905.11481.

[6] H. Tang, D. Key, and K. Ellis. *WorldCoder, a Model-Based LLM Agent: Building World Models by Writing Code and Interacting with the Environment.* 2024. arXiv: 2402.12275.

[7] NIST/SEMATECH. *e-Handbook of Statistical Methods: What is Acceptance Sampling?* URL: https://www.itl.nist.gov/div898/handbook/pmc/section2/pmc21.htm (visited on 09/05/2026).

[8] NIST/SEMATECH. *e-Handbook of Statistical Methods: Choosing a Sampling Plan with a Given OC Curve.* URL: https://www.itl.nist.gov/div898/handbook/pmc/section2/pmc222.htm (visited on 09/05/2026).

[9] NASA. *State-of-the-Art Small Spacecraft Technology: Thermal Control.* URL: https://www.nasa.gov/smallsat-institute/sst-soa/thermal-control/ (visited on 09/05/2026).

[10] J. Schnakenberg. « Network theory of microscopic and macroscopic behavior of master equation systems ». In: *Reviews of Modern Physics* 48 (1976), pp. 571–585. DOI: 10.1103/RevModPhys.48.571.

[11] A. van der Schaft and D. Jeltsema. « Port-Hamiltonian Systems Theory: An Introductory Overview ». In: *Foundations and Trends in Systems and Control* 1.2–3 (2014), pp. 173–378. DOI: 10.1561/2600000002.

[12] A. J. van der Schaft and B. M. Maschke. « Port-Hamiltonian Systems on Graphs ». In: *SIAM Journal on Control and Optimization* 51.2 (2013), pp. 906–937. DOI: 10.1137/110840091.

[13] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada. « Control Barrier Function Based Quadratic Programs for Safety Critical Systems ». In: *IEEE Transactions on Automatic Control* 62.8 (2017), pp. 3861–3876. DOI: 10.1109/TAC.2016.2638961. arXiv: 1609.06408.

[14] M. Dresher. *A Sampling Inspection Problem in Arms Control Agreements: A Game-Theoretic Analysis.* Tech. rep. RM-2972-ARPA. RAND Corporation, 1962. URL: https://www.rand.org/pubs/research_memoranda/RM2972.html.

[15] T. Gneiting and A. E. Raftery. « Strictly Proper Scoring Rules, Prediction, and Estimation ». In: *Journal of the American Statistical Association* 102.477 (2007), pp. 359–378. DOI: 10.1198/016214506000001437.

[16] D. Monderer and L. S. Shapley. « Potential Games ». In: *Games and Economic Behavior* 14.1 (1996), pp. 124–143. DOI: 10.1006/game.1996.0044.

[17] L. E. Blume. « The Statistical Mechanics of Strategic Interaction ». In: *Games and Economic Behavior* 5.3 (1993), pp. 387–424. DOI: 10.1006/game.1993.1023.

[18] V. Auletta, D. Ferraioli, F. Pasquale, and G. Persiano. *Mixing Time and Stationary Expected Social Welfare of Logit Dynamics.* 2010. arXiv: 1002.3474.

[19] S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon. « Time-uniform, nonparametric, nonasymptotic confidence sequences ». In: *The Annals of Statistics* 49.2 (2021), pp. 1055–1080. DOI: 10.1214/20-AOS1991. arXiv: 1810.08240.

[20] A. Ramdas, P. Grunwald, V. Vovk, and G. Shafer. *Game-theoretic statistics and safe anytime-valid inference.* 2023. arXiv: 2210.01948.

[21] National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* Tech. rep. NIST AI 100-1. NIST, 2023. DOI: 10.6028/NIST.AI.100-1.

[22] A. Prsa, P. Harmanec, G. Torres, et al. « Nominal Values for Selected Solar and Planetary Quantities: IAU 2015 Resolution B3 ». In: *The Astronomical Journal* 152.2 (2016), p. 41. DOI: 10.3847/0004-6256/152/2/41. arXiv: 1605.09788.

[23] NASA. *Automated Reconfigurable Mission Adaptive Digital Assembly Systems (ARMADAS).* External technology context; no integration or affiliation claimed. URL: https://www.nasa.gov/automated-reconfigurable-mission-adaptive-digital-assembly-systems-armadas/ (visited on 09/05/2026).