MONTREAL.AI · QUEBEC.AI

RESEARCH & USE EDITION / KNOWLEDGE, REPAIR, CONTINUITY

SUCCESSOR Ω — AEON Ω

The Continuity Estate

Repairable intelligence. Cooperative infrastructure.
A practical path from research to qualified use.

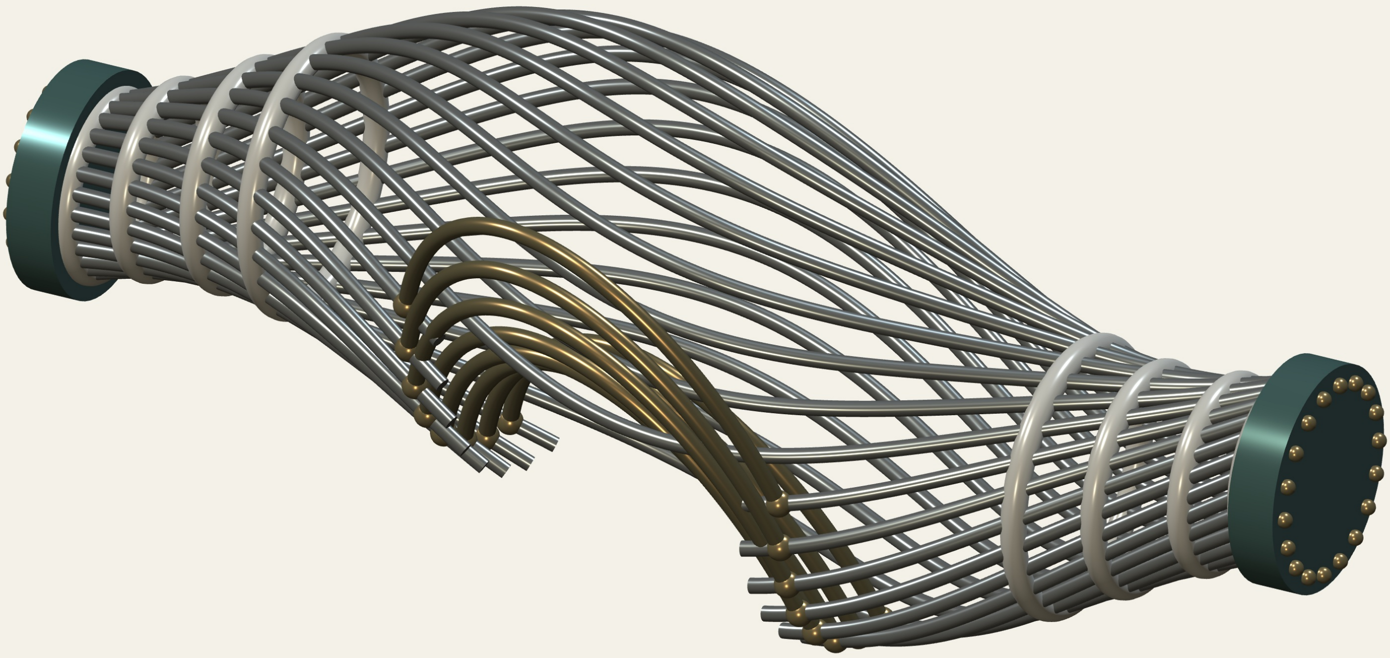

CONCEPTUAL COMPUTATIONAL ART / NOT AN APPARATUS OR ENGINEERING DRAWING

Vincent Boucher

President, MONTREAL.AI & QUEBEC.AI

AEON-TR-01 / 1.1

6 September 2026

Publication record and evidence contract

Title	SUCCESSOR Ω — AEON Ω
Subtitle	The Continuity Estate: Repairable intelligence, cooperative infrastructure and the price of reliable scale
Author	Vincent Boucher, MONTREAL.AI & QUEBEC.AI
Research/use edition	AEON-TR-01 / 1.1
Product lineage	NEURAL-SYMBOLIC SUCCESSOR Ω 10.2.0, unchanged
Empirical basis	AEON simulation 1.0; eight synthetic cohorts
New contribution	Finite repair resources, dependence-aware evidence, robust coalition stability
Status	AI-assisted research manuscript; not independently peer reviewed

The proposition

Own the capacity to preserve useful function when the equipment, world model, supplier and operating environment change. Continuity is not an inherited claim: it is maintained service supported by current evidence.

The retained simulation is the factual basis for this monograph. The new repair, dependence and coalition analyses are publication-side mathematical experiments; they do not change the first-party product or retroactively improve its original trial. Derivations are specialized consequences of established methods, not claims of previously unknown mathematics.

Executed	Local training, symbolic revision, finite planning, numerical reconstruction and reproducibility checks.
Derived	Conditional mathematical statements with stated assumptions and proofs.
Assumed	The fictional owner, rates, topology, prices, hardening effects, observation channels and investment horizons.
Prospective	Real customer Alpha, independently admitted Specialist ASI, factory control and stellar infrastructure.

No customer receipt, independent attestation, author signature, equipment installation or production permission is fabricated. Original product licences remain applicable. Conceptual images are not measurements; numerical figures are computed from disclosed data and equations.

Abstract

SUCCESSOR Ω—AEON Ω studies an institution that learns how components behave, plans for the service that interdependent assets deliver, and preserves evidence through successive models. A retained customer extension of the unchanged v10.2 scientific engine trains eight neural guides and 24 symbolic generations. All eight G2 candidates pass their nominal predictive protocol, all eight fail a later authored shift, and eight G3 candidates recover local validity. The 768-portfolio primary comparison produces a CA$1.079 billion gross authored annual contrast over risk ranking, but loses CA$8.617 million after additional costs to selected generic Beta and CA$9.001 million to an informed conventional comparator. Useful model-based decisions therefore do not establish a proprietary premium. The retained cooperative study has 57 nonempty cores among 64 games. A publication-side robust least-core analysis reduces that count to 29 under a declared 10% envelope; the relaxation is a dissatisfaction bound, not newly created money. Two retained grand coalitions have negative surplus, exposing the importance of an explicit decentralized fallback. A separate finite-source repair model examines a 40-unit cluster. Four crews satisfy a 95% service criterion under independent failures. None of the 12 tested staffing levels satisfies it under correlated environmental stress; a priced, hypothetical hardening intervention restores feasibility with four crews. Stationary equations, conservation residuals and relative-entropy contraction are checked. The original reserve-project campaign is preserved at CA$12.063 million conditional lower-envelope NPV. A new dependence analysis shows why marginal uptime does not prove reserve availability during an outage. Paired synthetic observations directly estimate that joint event; a separately priced campaign supports CA$18.323 million lower-envelope project NPV under fixed economic assumptions. This is not a financial-return confidence guarantee or a full-product Alpha claim. The formal treatment connects reliability importance, repair queues, non-equilibrium master equations, port-Hamiltonian physical networks, coalition stability, proper scoring, sequential claim budgets and finite-speed control. The long horizon is a repairable scientific-industrial estate, with useful power, heat, repair throughput and local autonomy physically explicit. No unrestricted discovery, independent Specialist ASI, field Alpha or production authority is asserted.

Keywords: neural-symbolic models; reliability; finite-source repair; common-mode risk; cooperative games; joint evidence; continuity; mission Alpha.

Owner decision brief

The strongest useful next mission

Qualify the service network, not just its component predictor. Measure whether a reserve works when the shared system fails, whether repair capacity remains available during stress, and whether the proposed coalition can retain its members without fictitious transfers.

Decision-relevant finding	Result
Current annual net comparison with informed Beta	CA$ -9.001 M
Original coalition cores / games	57 / 64
Strict robust cores at 10% envelope	29 / 64
Four-crew service, nominal / environmental stress	97.345% / 90.900%
Four-crew service with priced hardening	95.183%
Joint-event campaign lower project NPV	CA$ 18.323 M
Credited real customer Alpha	CA$ 0

The repair cluster, coalition envelopes and new paired-event campaign are separate publication-side analyses. Their amounts are not added to one another or to the original maintenance contrast. The owner keeps useful models and substitutes Beta where it is better. A new product premium needs new evidence.

Three reading routes

Owner: mission, comparative economics, reserve evidence, qualification and the continuity portfolio. **Scientist**: learned hypotheses, common dependence, repair generators, entropy and derivations. **Engineer/verifier**: source map, frozen artifacts, intervention feasibility, stationary and coalition witnesses, adverse cases and reproduction commands.

Contents

I / THE OWNED MISSION

Continuity is a productive asset

1 The institution that preserves useful service

Dr. Leïla Armand, a fictional control physicist, directs a CA$480 billion industrial estate across 48 campuses. Manufacturing, computation, scientific instruments and shared infrastructure depend on one another. The valuable object is not just installed capacity: it is useful service delivered when components age, crews are scarce and an upstream dependency fails. Her mandate is to preserve the service rather than a preferred model, price each intervention, inspect the strongest alternative and retain the evidence when the answer changes. The optimal owner programme is conditional on a mission, an information set and feasible alternatives. It is not a universal investment recommendation. Armand may retain a conventional predictor, fund a measurement campaign, buy spare capacity or defer construction. Her institutional advantage would lie in making and revisiting those decisions credibly, not in requiring the in-house candidate to win. The architecture explicitly makes cognitive suppliers replaceable while retaining the mission, programs, policies and Chronicle.[1]

The mission constitution

Minimize expected shortfall cost plus paid intervention cost for four-campus portfolios. Respect crew, spare-family, intervention-count and cash constraints. Predict in a declared normalized domain. Activate no equipment, procurement, payment or final admission connector.

The operating ambition then expands in a specific order: learn component behaviour; reason over supplied dependencies; establish the economics of intervention; negotiate shared benefits; measure reserve performance; and qualify repairable physical capacity. Each step has a separate counterfactual. Earlier manufacturing or design studies contribute context, not inherited financial gains.

2 What v10.2 executes, and what the customer adds

The retained AEON delivery contains a 55-file first-party snapshot of NEURAL-SYMBOLIC SUCCESSOR Ω 10.2.0. Those files are unchanged. Native scientific routines generate candidate expression programs, fit coefficients, train a small neural surrogate, propose an observation and evaluate frozen theories. The numerical problem, diagnostic observations, network topology and economic assumptions are customer-side inputs.[2, 3]

Layer	Implemented in this study	Not established by that fact
Native science	Search, symbolic evaluation, fitting, neural guide, ancestry	General causal discovery or arbitrary scientific theory invention
Customer adapter	Topology-aware intervention planning and explicit ledgers	A qualified plant controller or simulator of all physical dynamics
Secondary studies	Coalitions, measurement bounds, reserve investments	Executed contracts, independent evidence custody or actual receipts
This monograph	Repair queues, dependence analysis, robust least-core, proofs	A new production product version

The core experimental path uses no external model provider. A 16-hidden-unit neural guide with 65 parameters is trained locally for 1,400 steps in each cohort. The customer extension additionally uses NumPy, SciPy and scikit-learn. They are separately installed dependencies, not customer-owned intellectual property. The product's full institutional admission service is not invoked by this scientific demonstration.

3 Experimental protocol and information boundaries

The refined primary protocol uses seeds 9100–9107, selected after a retained exploratory run at 8100–8107 exposed insufficient search. This is a disclosed development revision, not externally sequestered confirmation. The authored generating family is deliberately compatible with the native expression language. Noise is uniform with half-width 0.0006; the declared maximum predictive error is 0.0015.

Stage, per cohort	Observations	Role
Initial narrow training / selection	36 / 24	Build G1 and neural guide
Nominal development / selection	129 / 64	128 new rows plus one proposed probe
Nominal held-out / stress	96 / 64	Evaluate G2 after freeze
Shift development / selection	160 / 64	Fit a descendant G3
Shift held-out / stress	96 / 64	Falsify G2 and evaluate G3

Each cohort also has 24 development portfolios for decision-level comparator selection, 96 nominal final portfolios and 32 shifted portfolios. Twelve component-level generic candidates are not claimed: the source fits ten generic regressors plus one informed model per cohort, giving 88 Beta fits. The primary study records 25,270 native candidate fits, 4,608 nominal decisions, 1,024 shifted decisions and 5,760 ledger checks. These are source counts, not measures of independent evidential strength. All observations remain under the same author's control. Distinct identifiers and frozen artifacts reduce accidental leakage, but do not create institutional independence. The bootstrap resamples eight cohort mean differences, not thousands of plans as if each were an independently trained system. The source's risk-ranking incumbent is also weaker than its informed conventional reference; both must remain visible.

4 Three generations of executable understanding

The model predicts an interval's unavailability from normalized age $a \in [0, 2]$ and load $l \in [0, 2.5]$. It does not learn time-to-failure distributions from censored observations. For cohort 9100, the retained initial expression is

$$\hat{q}_1(a, l) = 0.041539532a^2 + 0.0099562412e^l. \tag{4.1}$$

The neural contribution changes the selected probe in all eight cohorts; the first cohort's contradiction is 0.0645137 in unavailability. The observation affects a later program, rather than being attached to an unchanged recommendation. Its G2 expression is

$$\hat{q}_2(a,l) = 0.074117477 + 0.0007205914l + 0.03512249a^2 + 0.015769323al - 0.064398403\cos l. \tag{4.2}$$

After the authored regime change, G2 fails the fixed protocol. G3 retains its ancestry but receives a new identity:

$$\hat{q}_3(a,l) = 0.074933377 + 0.034973548a^2 + 0.016029658al + 0.025026078l^2 - 0.064945882\cos l. \tag{4.3}$$

The structural l^2 term is the visible successor change. It approximates a known authored perturbation, not a newly discovered law. A data-only typed expression and declared input envelope support inspection. They do not guarantee that the abstraction is causal, complete or safe outside its observed domain. WorldCoder provides a related research direction for executable world models; its results are not inherited as evidence for AEON.[4]

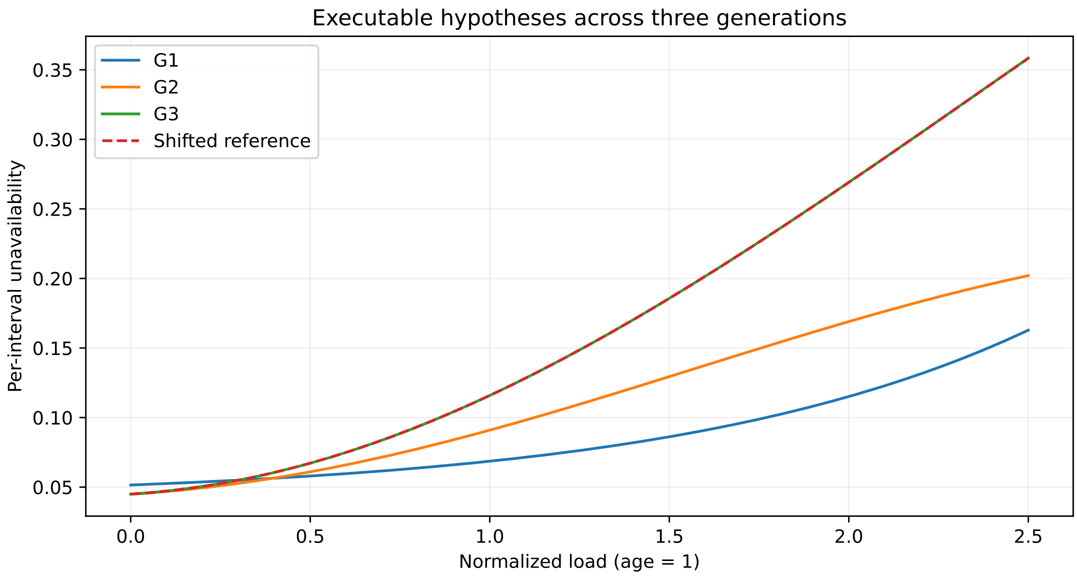

Figure 1. Retained cohort-zero predictions across three generations. Raw model behaviour and the shifted reference remain visible; the ordinate is interval unavailability, not a hazard rate.

5 Dependency-aware decisions and their error budget

For binary intervention vector $x \in \{0,1\}^{10}$, let d_i denote planned unavailability and r_i residual post-intervention unavailability. The source defines

$$v_i(x_i) = (1-x_i)(1-\hat{q}_i) + x_i(1-d_i)(1-r_i). \tag{5.1}$$

Four local parallel pairs feed through shared components 9 and 10. Under the nominal independence assumptions,

$$A_j(x) = [1-(1-v_{2j-1})(1-v_{2j})]v_9v_{10}(1-q_c), \quad j = 1,\dots,4. \tag{5.2}$$

The objective is

$$J(x) = \sum_{j=1}^{4} D_j m_j[1-A_j(x)] + \sum_{i=1}^{10} c_i x_i, \quad \min_{x\in\mathcal{X}} J(x), \tag{5.3}$$

where D_j is MWh of demanded service, m_j its opportunity cost in CAD/MWh, and $\mathcal{X}$ enforces intervention count, 32 crew-hours, spare families and CA$115,000 cash. The logical topology is supplied; it is not a power-flow calculation or a topology learned autonomously from telemetry.

Proposition 5.1 (Importance is a property of the system). For a local component i of campus j, $\partial A_j/\partial v_i$ contains the partner's unavailability. For shared component 9, its marginal contribution sums across all campuses:

$$-\frac{\partial J}{\partial v_9} = (1-q_c)v_{10}\sum_j D_j m_j[1-(1-v_{2j-1})(1-v_{2j})]. \tag{5.4}$$

Ranking q_i alone is therefore not a ranking of intervention value.

This is direct differentiation of the reliability polynomial; the independence condition is indispensable.[5] Shared series dependencies create complementarity: restoring one has less value while another remains unavailable. Local parallel redundancy creates a different interaction. A universal submodularity claim, or a guaranteed greedy approximation based on it, is not justified.

Proposition 5.2 (Probability error to decision regret). For feasible x, suppose $|q_i - \hat{q}_i| \le \epsilon_i$, costs and topology are correct, and let $I_j = \{2j-1, 2j, 9, 10\}$. Then

$$|J(x) - \widehat{J}(x)| \le \Delta := \sum_j D_j m_j \sum_{i\in I_j} \epsilon_i. \tag{5.5}$$

If $\hat{x}$ minimizes the modeled objective over the same feasible set and x^* minimizes truth, $J(\hat{x}) - J(x^*) \le 2\Delta$.

The bound follows by telescoping products in $[0,1]$ and inserting $\widehat{J}$ on each side of the optimizer inequality. It cannot absorb an unknown shared failure, an incorrect action effect or a different feasible set. Those change the model class rather than a coefficient error.

6 Results: preserve value without inventing a premium

The first final portfolio has 146 feasible intervention combinations. Risk ranking selects components 2, 5 and 6; network planning selects 1, 5 and 10. Expected service increases from 127,602 to 128,108 MWh, while intervention cash falls from CA$104,328 to CA$60,243. Total expected cost falls by CA$172,248. These are conditional expectations for one synthetic interval, not realized production.

Primary comparison, authored annual scale	Mean net, M CAD	Cohort interval, M CAD
Risk-ranked incumbent	1 069.966	[1 026.358, 1 108.766]
Selected generic Beta	-8.617	[-9.001, -7.861]
Informed conventional Beta	-9.001	[-9.003, -9.000]
Informed greedy planner	-7.084	[-7.344, -6.834]

Annualization multiplies four-campus interval differences by $12 \times 52 = 624$, then subtracts CA$9 million of assumed incremental cost. The positive incumbent contrast and negative Beta comparisons are not additive. No selected or informed Beta cohort has positive net Alpha. The annualization presumes transferable mean effects and does not diversify correlated estate-wide losses.

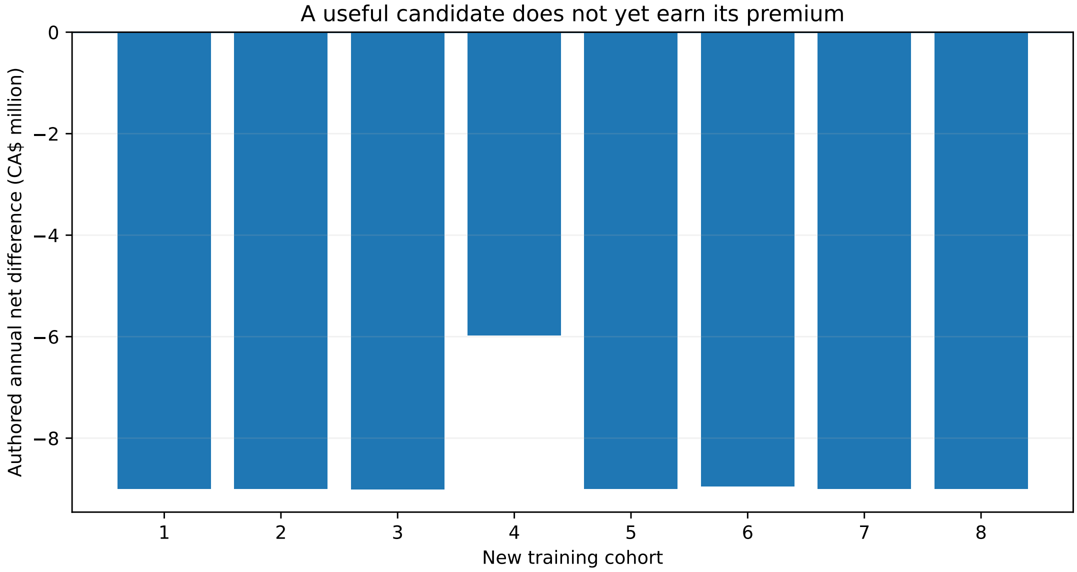

Figure 2. Retained cohort-level competitive results. The relevant baseline is not restricted to the weakest incumbent.

The informed comparator is a conventional least-squares model supplied with the authored mechanism's features and fitted on the same noisy observations. It is not an oracle receiving the final labels, but it is a deliberately informed diagnostic alternative, not a verified vendor quotation. Its near-identical choices show why a useful explanatory program can still fail to earn an exclusive-model premium. The customer should preserve the method and use the best suitable component.

II / PHYSICAL CONTINUITY

Repair is part of the system

7 An unavailability probability is not a failure clock

The native AEON program estimates a probability over one decision interval. It does not identify a unique continuous-time failure intensity or repair distribution. Even in a two-state exponential model, stationary unavailability $q = \lambda/(\lambda+\mu)$ does not identify either rate: $(s\lambda, s\mu)$ has the same q for every $s > 0$, but a different turnover rate and transient response. A probability accurate for a weekly decision cannot be silently treated as an annual hazard.

Proposition 7.1 (Availability non-identifiability). For a two-state unit with up-to-down rate λ and repair rate μ, stationary availability is $a = \mu/(\lambda+\mu)$ and failure-start throughput is $\nu = \lambda\mu/(\lambda+\mu)$. A common positive rate scaling preserves a and multiplies ν by that scale.

The proof is the balance equation $(1-a)\mu = a\lambda$. The implication is practical: identical reported uptime can require radically different spare inventories and repair work. The new queue study therefore declares its rates independently; it does not infer them from the retained learned expressions. General repair modelling also needs age, censoring, competing failures, maintenance selection, repair quality and common exposures. NASA's PCoE repository is a possible future data resource, not data used here.[6]

8 Finite-source repair: the hidden capacity constraint

Consider N statistically identical modules and c repair crews. Let $K(t)$ count failed modules. Working modules fail independently at rate λ per day; each occupied crew repairs at rate μ per day. The new calculation assumes exponential durations, immediate crew assignment, unlimited compatible spares and no travel time. It belongs to the classical finite-source machine-repair family.[7, 8]

$$Q_{k,k+1} = (N-k)\lambda, \qquad Q_{k,k-1} = \min(k,c)\mu, \qquad Q_{kk} = -\sum_{j\neq k} Q_{kj}. \tag{8.1}$$

Proposition 8.1 (Repairable stationary capacity). For $1 \le c \le N$ and strictly positive rates, the unique stationary distribution is

$$\pi_k = \frac{w_k}{\sum_{j=0}^{N} w_j}, \qquad w_0 = 1, \quad w_k = \prod_{i=1}^{k} \frac{(N-i+1)\lambda}{\min(i,c)\mu}. \tag{8.2}$$

Mean module availability is $a_c = 1 - \mathbb{E}_\pi[K]/N$. For a requirement of at least m working modules, the stationary service probability is $\sum_{k=0}^{N-m} \pi_k$.

Detailed balance gives the recurrence; positivity and finite irreducibility give uniqueness. At $c = N$, the failed count is binomial with probability $\lambda/(\lambda+\mu)$. For finite crews, that independent-unit shortcut need not hold. A finite closed repair system always has a stationary distribution under these assumptions, even when $N\lambda > c\mu$; the penalty is reduced availability, not an infinite population queue exploding without bound.

Proposition 8.2 (A necessary staffing inequality). In the same model, $\lambda N a_c = \mu \mathbb{E}[\min(K, c)] \le c\mu$. Therefore achieving mean availability at least a_* requires $c \ge N\lambda a_*/\mu$. This is necessary, not sufficient for either mean or tail service.

The flow identity follows by applying stationarity to K. It is the missing accounting line in an estate described only by uptime targets: failed units must enter repair at the same long-run rate at which completed repairs return them. A safety stock does not eliminate this work.

9 A common environment changes the repair decision

The new cluster has 40 units of 25 MW each, $\lambda = 0.002$/day and $\mu = 0.05$/day. Its criterion is at least 36 units delivering service with 95% stationary probability. The objective prices expected lost output at CA$150/MWh and each crew at CA$1.8 million/year. These are authored values, not market data. Twelve crew choices are evaluated exactly within the finite model. Under independent failures, four crews minimize the declared cost among feasible options. Service probability is 97.345%, at a modeled annual total cost of CA$59.013 million. Add a binary environment $E(t)$: normal-to-stress rate 0.002/day, stress-to-normal 0.098/day. Stress occupies 2% of time, multiplies the failure intensity by eight, halves repair intensity, and blocks delivered service. A joint chain over (E, K) has 82 states. It preserves backlog after the environment recovers.

$$A_{\text{service}}(c) = \sum_{k=0}^{N-m} \pi_{0,k}, \qquad \bar{a}_{\text{delivered}}(c) = \frac{1}{N}\sum_{k=0}^{N}(N-k)\pi_{0,k}. \tag{9.1}$$

Four crews now achieve only 90.900% service probability. None of the 12 staffing options attains 95%; twelve crews reach approximately 94.313%. Simply multiplying nominal service by 98% misses the persistent backlog. A known fraction of stressed time is not enough to justify an independence approximation. A separate

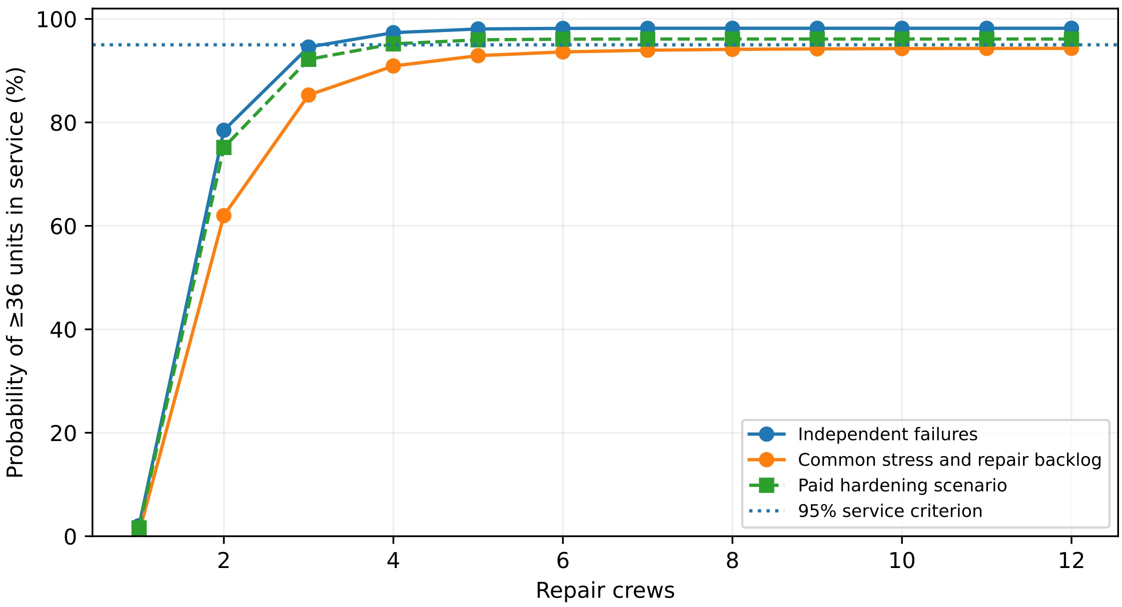

Figure 3. New finite-cluster calculation. The same staffing choice can pass an independent model and fail a joint environmental model.

hypothetical hardening option costs CA$3 million annually, limits the stress failure multiplier to two and preserves normal repair speed. It does not remove the common outage. With this paid option, four crews reach 95.183% and cost CA$88.037 million/year in the authored account. The unstressed model's cheaper figure is not an achievable comparison in the same world. Hardening effects require physical evidence before real procurement. Fifteen of 27 sensitivity settings have no acceptable design among the 12 tested staffing levels.

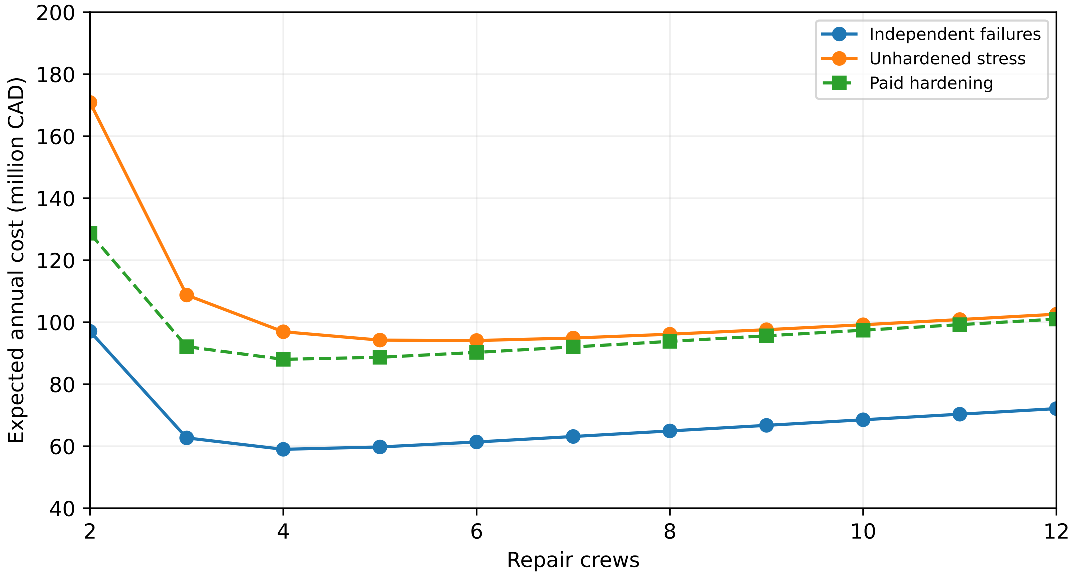

Figure 4. Expected lost-margin and crew cost, with hardening charged separately. The vertical range focuses on the useful staffing region and excludes the one-crew extreme.

10 Master equations and what entropy does not prove

Let a row probability vector evolve by $\dot{p} = pQ$. This is a probabilistic master equation. Its transition generator encodes the assumed failure and repair mechanism, not a fundamental Hamiltonian of the estate. Network representations and thermodynamic analogies are useful only when their semantics are made explicit.[9, 10]

Proposition 10.1 (Relative-entropy contraction). For a finite irreducible generator with stationary $\pi_i > 0$ and positive p_i, let $r_i = p_i/\pi_i$. Then

$$\frac{\mathrm{d}}{\mathrm{d}t} D(p\|\pi) = \sum_{i\neq j} \pi_i r_i Q_{ij} \log \frac{r_j}{r_i} \leq 0. \tag{10.1}$$

Use $\log x \leq x - 1$ and stationarity to bound the right side by $\sum_{ij} \pi_i Q_{ij}(r_j - r_i) = 0$. No detailed-balance assumption is required. This says the distribution approaches its stationary law. It does not say the law has adequate service. A system can converge stably to unacceptable availability. The transient calculation starts near 12 failed units, records nine times through day 160, and verifies nonpositive derivatives. The distinct stationary cycle-production expression is $\frac{1}{2}\sum_{ij}(\pi_i Q_{ij} - \pi_j Q_{ji}) \log[\pi_i Q_{ij}/(\pi_j Q_{ji})] \geq 0$ for reversible edges. It can remain positive in a non-equilibrium stationary state while $D(p\|\pi) = 0$. Converting that information-rate expression into physical heat requires local detailed balance and a justified mapping of states to energies and reservoirs. The maintenance chain alone does not supply that mapping.

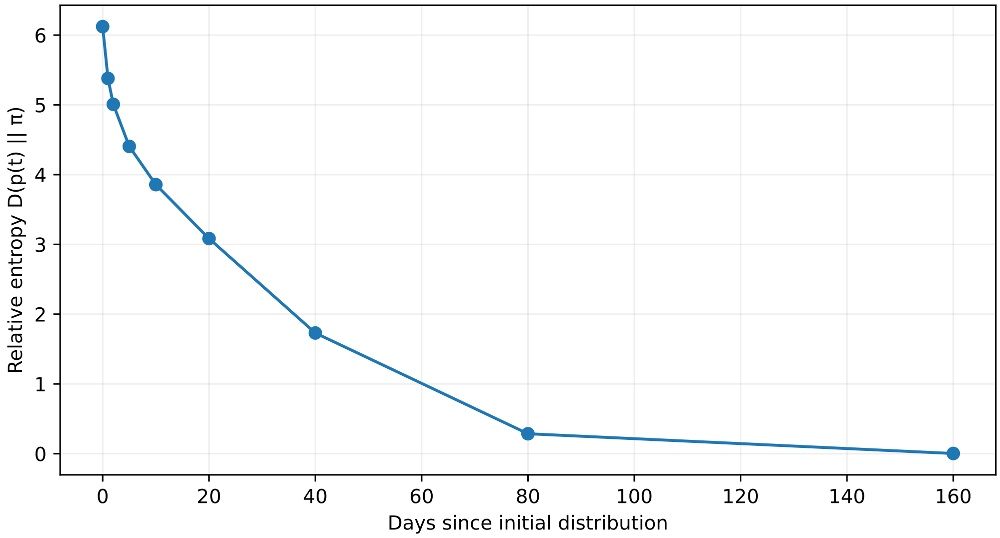

Figure 5. Relative entropy of the new joint repair model approaches its stationary law. This is a numerical check of probabilistic relaxation, not a safety certificate.

11 Physical network control beyond logical availability

An availability block is not a model of voltages, flow, temperature or mechanical energy. A physical adapter can instead expose a port-Hamiltonian structure, with incidence-derived interconnection, state z, storage H and ports (u, y):[11, 12]

$$\dot{z} = [J(z) - R(z)]\nabla H(z) + G(z)u, \quad J = -J^\top,\ R \succeq 0,\ y = G^\top \nabla H. \tag{11.1}$$

Proposition 11.1 (Power and dissipation).

$$\dot{H} = -\nabla H^\top R \nabla H + y^\top u. \tag{11.2}$$

A switching repair configuration must preserve its declared interconnection constraints and account for reset energy; repairing a logical node does not waive this balance.

For a thermal network with heat capacities $C_i > 0$, symmetric conductances $g_{ij} \geq 0$ and absolute temperatures $T_i > 0$, define availability relative to bath T_0 by

$$\mathcal{F}(T) = \sum_i C_i \left[T_i - T_0 - T_0 \log(T_i/T_0)\right]. \tag{11.3}$$

Proposition 11.2 (Conductive exergy loss). For $C_i\dot{T}_i = \sum_j g_{ij}(T_j - T_i)$,

$$\dot{\mathcal{F}} = -T_0 \sum_{i<j} g_{ij} \frac{(T_i - T_j)^2}{T_i T_j} \leq 0. \tag{11.4}$$

External heat generation, radiation and cooling add their own port terms. The identity helps detect a physically impossible repair or routing proposal; it does not identify the correct conductances. Six separately seeded numerical networks verify the algebraic identity in the publication tests.

Proposition 11.3 (A sufficient robust thermal barrier). For $C\dot{T} = P(t) - Q(T, u, \theta)$ and $h = T_{\max} - T$, assume a lower certified heat-removal bound $Q \geq Q_{\min}$ and input bound $P \leq P_{\max}$. Enforcing

$$Q_{\min}(T, u) \geq P_{\max}(t) - C\alpha(T_{\max} - T), \quad \alpha > 0, \tag{11.5}$$

gives $\dot{h} \geq -\alpha h$ and preserves $T \leq T_{\max}$ from a safe initial state, subject to a feasible input and the assumed dynamics.

This is a comparison argument within the control-barrier framework.[13] Sampling delay, actuator saturation, uncertain thermal mass and failed sensors must be included before implementation. None of these equations is represented as a connected AEON plant controller.

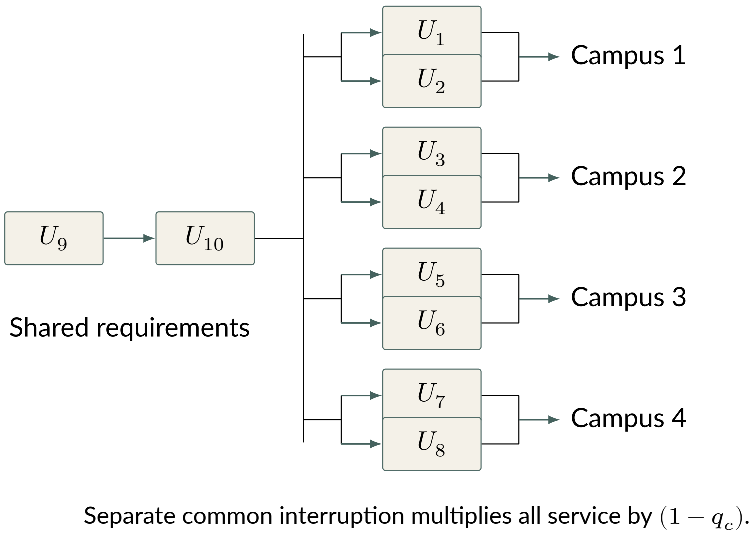

Figure 6. Logical reliability structure supplied to the retained study. Each local pair is parallel; shared components are required in series. This is not an electrical, hydraulic or thermal flow network.

III / COOPERATIVE CONTINUITY

Shared value must survive disagreement

12 An efficient network is not automatically a stable institution

AEON's second problem is collective: four campuses can combine resources, finance shared dependencies and allocate the resulting cost reduction. Let $C(S)$ be the minimum modeled cost for coalition S, and $C_i = C(\{i\})$. The characteristic surplus is

$$v(S) = \sum_{i \in S} C_i - C(S), \qquad v(\emptyset) = 0. \tag{12.1}$$

The source evaluates all 16 subsets for 64 newly generated portfolios. Its mean grand-coalition surplus is CA$2.238 million per four-campus interval. This comparison changes resource pooling and therefore answers a different question from model superiority. It is not added to the primary annual contrast. The surplus assumes full information, transferable utility and no negotiation or enforcement cost. The Shapley allocation is the average marginal contribution across joining orders:[14]

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n-|S|-1)!}{n!}[v(S \cup \{i\}) - v(S)]. \tag{12.2}$$

Its efficiency $\sum_i \phi_i = v(N)$ is not a guarantee that every subgroup prefers to remain. The core requires

$$\mathcal{C}(v) = \{x : \sum_i x_i = v(N), \quad \sum_{i \in S} x_i \geq v(S) \quad \forall S \subset N\}. \tag{12.3}$$

Seventeen retained Shapley allocations violate a coalition condition. Fifty-seven games have a feasible core; seven do not. The platform must distinguish "allocated the total" from "passed the declared stability conditions." Existence results for cooperative cores depend on their assumptions; nonemptiness is not guaranteed for an arbitrary industrial game.[15]

13 Least-core certificates and robust participation

A least-core calculation identifies the smallest uniform relaxation $\epsilon \geq 0$ that makes the subgroup constraints jointly feasible:[16]

$$\begin{aligned} &\min_{x, \epsilon \geq 0} \ \epsilon \\ &\text{s.t.} \ \sum_i x_i = v(N), \quad \sum_{i \in S} x_i + \epsilon \geq v(S) \quad (\emptyset \neq S \subsetneq N). \end{aligned} \tag{13.1}$$

We truncate any negative classical least-core radius at zero. Allocations x_i are unrestricted in this diagnostic optimization; singleton constraints are relaxed too. A negative allocation or positive ϵ is not an individually rational contract. The result quantifies modeled dissatisfaction. It is neither a subsidy automatically paid nor a stable agreement.

Proposition 13.1 (A dual certificate of unresolved coalition pressure). For nonnegative weights α_S and a common β satisfying $\sum_{S\ni i}\alpha_S = \beta$ for every player and $\sum_S \alpha_S \le 1$, every feasible solution of (13.1) obeys

$$\epsilon \ge \sum_S \alpha_S v(S) - \beta v(N). \tag{13.2}$$

The optimal such lower bound equals the primal optimum when the finite linear programmes are feasible.

Multiply and sum coalition constraints; each x_i appears with coefficient β, then use the budget equality. Strong linear-programming duality gives equality at optimum. A certificate names the coalition combinations causing conflict rather than concealing them behind a scalar score. For the new robustness study, use $\underline{v}(N) = v(N) - \eta|v(N)|$ and $\overline{v}(S) = v(S) + \eta|v(S)|$. Require a baseline allocation totaling the guaranteed grand surplus and satisfying the upper subgroup claims. If actual grand surplus is higher, nonnegative residual distributions can complete the allocation. This is a rectangular scenario envelope, not a joint confidence region and not a new physically simulated game.

Envelope η	0%	2%	5%	10%
Strict feasible allocations / 64	57	46	38	29
Mean relaxation, CAD	780	7 438	29 812	86 552
Maximum relaxation, CAD	15 543	97 881	282 155	596 844

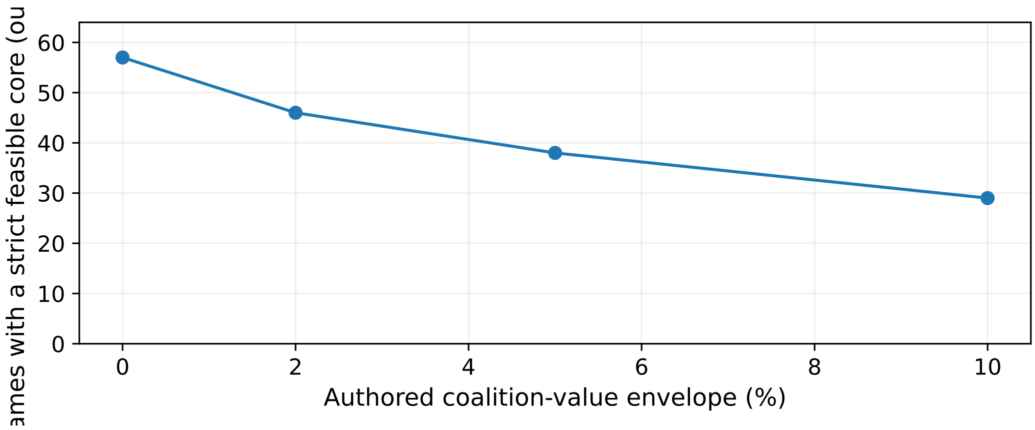

Figure 7. The number of feasible strict allocations contracts as declared subgroup uncertainty expands. This is a fixed-record sensitivity analysis.

Two original grand coalitions have negative surplus: approximately CA$2,487 and CA$13,402 below separate operation. The source is retained unchanged. Its joint action menu has a global three-intervention limit and shared inventories; independently optimized singleton menus need not fit that grand menu. Therefore negative surplus is possible. A real voluntary federation should explicitly permit a feasible decentralized fallback; that would define a different game and requires a new evaluation, not a silent numerical patch.

14 Truthful forecasts are not the same as diligent science

A forecast contract should not pay for a preferred conclusion. For a Bernoulli event Y with true conditional probability p, the quadratic score $S(r, Y) = 1 - (r - Y)^2$ has

$$\mathbb{E}[S(r, Y) \mid p] = 1 - p(1 - p) - (r - p)^2. \tag{14.1}$$

Truthful reporting uniquely maximizes expected score for a risk-neutral reporter who cannot manipulate the event, the evaluation population or payment eligibility. Those restrictions are essential.[17]

Proposition 14.1 (The information-acquisition premium). Let $p_0 = \mathbb{E}[Y \mid \mathcal{F}]$ and $p_1 = \mathbb{E}[Y \mid \mathcal{F}, Z]$. The expected improvement in an optimally reported quadratic score from observing Z is

$$\mathbb{E}[S(p_1, Y) - S(p_0, Y)] = \mathbb{E}[(p_1 - p_0)^2]. \tag{14.2}$$

A score multiplier b motivates the observation only if this expected reward exceeds its private cost and opportunity cost.

The identity follows from the orthogonality of conditional-expectation residuals. It distinguishes honesty from effort, but still rewards prediction rather than customer decision value. An assay can improve the score without changing any action. AEON therefore evaluates downstream service and costs separately and does not treat a proper score as an Alpha certificate. A second contract concerns repair effort. If diligent work costs $c_H - c_L$, raises accepted success probability by Δp, and a shirker faces incremental detected-and-enforced loss $adeB$, high effort is incentivized only when $\Delta p\, b + adeB \geq c_H - c_L$. This algebraic condition assumes the relevant probabilities and enforceable liability are known. A nominal bond is not effective liability if the counterparty cannot pay. False sanctions, common verifiers and hidden effort require further contract design.

15 Scheduling as a potential game, not a metaphor

A research federation may allocate finite repair slots by a congestion mechanism. Suppose each member's cost is a private work cost plus resource charges $\ell_e(n_e)$ for occupied resources. If incentives are exactly aligned, Rosenthal-type cumulative charges define

$$\Phi(a) = -\sum_i c_i(a_i) - \sum_e \sum_{k=1}^{n_e(a)} \ell_e(k). \tag{15.1}$$

A unilateral payoff change equals the change in Φ, so finite strict-improvement paths terminate at pure equilibria. Such equilibria are not necessarily globally efficient. The AEON coalition solver is not this decentralized mechanism; it is a separate full-information linear programme.[18]

Proposition 15.1 (Finite logit equilibrium). If a single player is selected with state-independent probability and updates with logit probabilities proportional to $\exp(\beta u_i)$ in a finite exact potential game, the stationary measure is

$$\pi_\beta(a) = Z^{-1} e^{\beta \Phi(a)}, \qquad \log Z = \sup_p \{\beta \mathbb{E}_p[\Phi] + H(p)\}. \tag{15.2}$$

The variational gap is $D(p \| \pi_\beta)$.

Transition ratios establish detailed balance for profiles differing at one coordinate; substitution establishes the variational identity.[19] Here $1/\beta$ is a decision-noise scale in utility units, not physical temperature. Simultaneous updates, state-dependent scheduling, infeasible intermediate profiles or nonpotential incentives may invalidate the formula. Stronger concentration on a potential maximum also need not imply fast mixing. This theory is included to specify an implementable coordination research target, not to rename arbitrary agent consensus as thermodynamics.

IV / EVIDENCE TO CAPITAL

Measure the event that pays

16 The retained reserve project

The source's separate investment problem concerns four campuses exposed to a common interruption. Binary design z_i installs an independent reserve at campus i. Its authored annual contribution is proportional to margin M_i, common-interruption fraction q and reserve availability r_i. Equipment costs, annual service, two installation slots, a CA$95 million budget, two-year commissioning and a twelve-year horizon are explicit. Year seven removes benefits but retains maintenance.

$$\mathrm{NPV}(z;\theta) = -K(z) - C_{\text{evidence}} + \sum_{t=1}^{12} \frac{\mathbf{1}\{t > 2\} \sum_i z_i [M_i \theta_i h_t - s_i]}{(1.08)^t}, \tag{16.1}$$

where $h_t = 0$ in the impaired year and otherwise declines by 6% per operating year. In the retained model $\theta_i = qr_i$. Before its measurement campaign, conservative values select no-build. Five channels, each with 4,096 synthetic Bernoulli exposures, produce simultaneous one-sided lower bounds by a 5% Bonferroni allocation. A CA$2.4 million campaign changes the lower-envelope optimum to reserves at campuses 1 and 2. The retained equipment bill is CA$63.348 million. After its campaign, nominal NPV is CA$23.622 million and lower-envelope NPV is CA$12.063 million. These results are not revised by the new paper. Exact binomial bounds are valid for the stated stationary Bernoulli parameters, not automatically for future project cash.[20]

17 The dependence question that marginal uptime cannot answer

A reserve is valuable specifically when the main service is interrupted. Let O be that interruption and R the event that reserve service is available. The revenue-relevant event is $O \cap R$, not a product of separately observed means. If $r = \Pr(R \mid O)$, then $\Pr(O \cap R) = qr$ by definition. If instead $r = \Pr(R)$ is marginal uptime, the product requires independence. The original simulation stipulates its channel model; the following study examines a distinct interpretation without silently correcting it.

Proposition 17.1 (Sharp joint-service bounds). Given only $q = \Pr(O)$ and $r = \Pr(R)$, the sharp bounds are

$$\max(0, q + r - 1) \le \Pr(O \cap R) \le \min(q, r). \tag{17.1}$$

Both endpoints are attained by nonnegative 2×2 probability tables with the given marginals.

The retained common fraction is 0.062274 and reserve marginal availabilities range from about 0.856 to 0.952. For three campuses the lower joint bound is zero; for the third campus it is about 0.014128. Under those conservative joint values, the same investment routine selects no-build. High marginal reliability does not prevent the reserve from failing exactly when the shared event occurs. This is not evidence that the retained independent model is empirically false: no field data are present. It identifies the measurement needed to justify its revenue factor. An independent power supply, common coolant loop, shared control software or correlated maintenance practice can each change the relevant conditional availability. Separate availability certificates are not a joint-service certificate.

18 Paired observations that change the capital decision

The new campaign observes $J_{ti} = \mathbf{1}\{O_t \cap R_{ti}\}$ directly at 8,192 common exposure times. Four sites therefore contribute 32,768 binary outcomes, but the four channels need not be independent. The shared outage draw is preserved across sites. Rows are assumed independent and representative over time. Data remain synthetic and under the same author's custody. The campaign costs CA$3.2 million; it does not reuse the old campaign's cost or its nominal confidence claim.

Proposition 18.1 (Joint evidence and simultaneous selection). Let $K_i = \sum_{t=1}^{n} J_{ti}$ and $\delta_i > 0$ with $\sum_i \delta_i \le \delta$. Define

$$L_i = \begin{cases} 0, & K_i = 0, \\ F^{-1}_{\mathrm{Beta}(K_i, n-K_i+1)}(\delta_i), & K_i > 0. \end{cases} \tag{18.1}$$

Under each channel's binomial sampling assumptions, $\Pr(\theta_i \ge L_i \; \forall i) \ge 1 - \delta$, regardless of dependence between channels. If all other project inputs are fixed and NPV is coordinatewise nondecreasing in θ, selecting a design after inspecting the simultaneous vector does not invalidate its lower bound with respect to these parameters.

The proof combines one-sided binomial inversion with a union bound, then monotonicity. It does not claim simultaneous coverage of uncertain prices, construction delays, future dependence, repair costs, taxes or demand. It is also not valid for arbitrary repeated optional stopping under the fixed-sample calculation.

Campus	1	2	3	4
Joint successes / 8,192	473	479	504	464
Simultaneous lower bound	0.052103	0.052801	0.055711	0.051056
Selected reserve	1	1	0	0

The new lower-envelope decision installs the first two reserves. Nominal NPV after this campaign is CA$22.822 million; lower-envelope project NPV is CA$18.323 million. Different observations, sample size and costs distinguish this study from the original CA$12.063 million record. Their difference is not a randomized estimate of the value of paired sampling. The complete product's CA$9 million annual premium is not included in either project-level calculation.

19 Proof budgets across successors

An institution that repeatedly manufactures candidates has repeated opportunities for false promotion. A single 95% interval copied across generations does not provide a 95% estate-wide statement. Nor may a common learned release count all its plans as independently trained replications.

Proposition 19.1 (A lifetime budget for false comparative claims). Before each fresh generation g, allocate conditional error budgets $\delta_{g,b}$ across comparators with $\sum_{g,b} \delta_{g,b} \le \delta$. Suppose each candidate is frozen before its fresh evidence and its lower bound is valid conditional on the previous history. Then the probability of any false comparative lower-bound claim is at most δ.

Conditional probability followed by the union bound proves the result. For example $\delta_g = \delta/[g(g+1)]$ is summable. Optional stopping requires a valid confidence sequence or an explicitly allocated inspection schedule, not repeated use of a fixed-time interval. Time-uniform methods provide an appropriate research foundation.[21] This mathematical budget does not protect against poisoned evidence, an unqualified instrument, a changing mission or verifier capture. The architecture's proof tuple binds mission, release, comparators, protocol, evidence, operating environment and validity window. Its diagram on page 21 separates

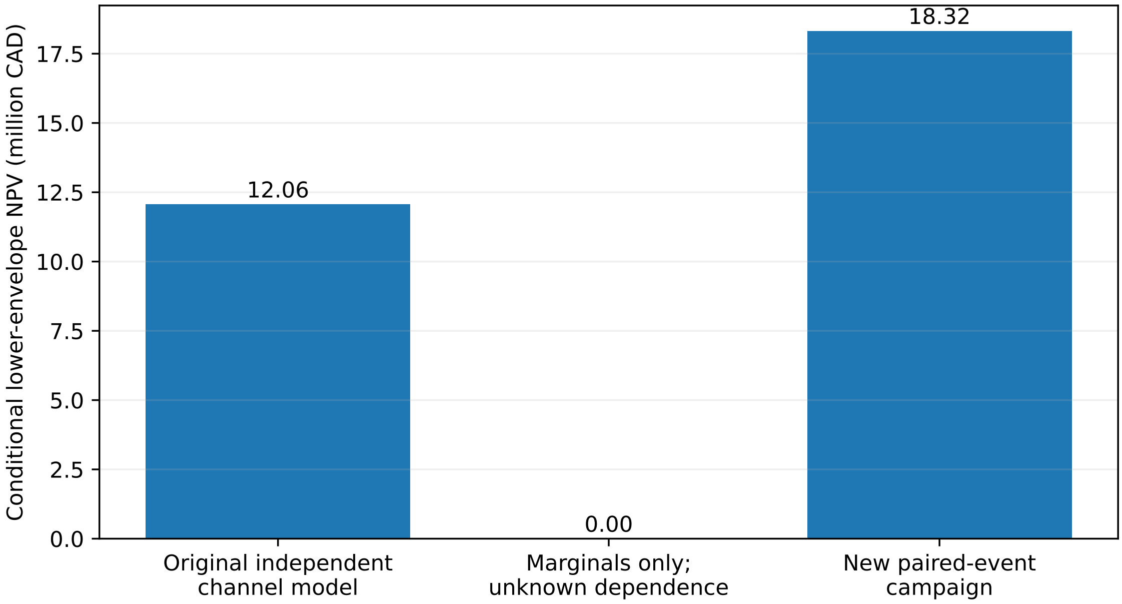

Figure 8. Three distinct information interpretations of a reserve project. The third bar uses a new, separately priced synthetic paired campaign; it is not a replacement for the original record.

the serving release from the learning branch; page 54 makes succession a new proof and admission event. AEON uses that doctrine, but its scientific archive does not execute the final admission service.[1] A defensible mission-specific Alpha certificate requires lower advantage against every relevant fixed comparator, less nonduplicated costs and reserves, above a positive admission threshold. A field violation of an immutable safety constraint is not tradable against a favorable mean return. In this publication the source's negative competitive result remains decisive: the candidate is not admitted as a Specialist ASI.

V / THE CONTINUITY ESTATE

Reliable scale must finance its own repair

20 From accepted service to available capital

In the retained twelve-year reserve path, initial construction and measurement require external opening capital; forecast gains cannot finance them before they arrive. The first two years have no operating benefit. In year seven, impairment eliminates benefits while maintenance continues. Positive operating net is split 25% to science, 25% to distributions and 50% to replacement reserves. Losses draw from those reserves.

$$C_{t+1} = C_t + R_t - O_t - I_t - V_t - D_t, \qquad C_t \geq C_{\min}, \tag{20.1}$$

where receipts R_t, operating expense O_t, investment I_t, verification/research V_t and distributions D_t have separate identities. The illustrative original path totals CA$148.036 million of operating net: CA$37.689 million each for research and distributions, and CA$72.658 million retained for replacement. These amounts are before subtracting the initial equipment and measurement expenditure. They do not constitute enterprise valuation, proof of repayment or financial independence of the project. A real owner must distinguish avoided opportunity loss, invoiced services, accepted deliveries, receivables, cash and retained earnings. The primary maintenance objective prices expected service loss. It does not create invoices or bank balances. The reserve account is a separately authored continuation. Cash and repair resources can be complements: sufficient money with no compatible spare or local crew does not restore service.

21 An owner programme that can earn stronger claims

The best use of present AEON is a bounded research, shadow-evaluation and evidence-design programme. Armand should first qualify the failure and repair measurements, then reconstruct real dependency topology and feasible intervention effects. She should compare the candidate with an available supplier and an informed conventional method under the same costs, operational constraints and information. A negative comparison is a valid result.

Gate	Decision	Required evidence
D0	Constitute one high-value mission	Owner, observable service, baseline, exclusions, research budget
D1	Build the measurement world	Rights, timestamps, calibration, censored histories, common events
D2	Freeze and challenge	Exact programs, independent custody, costs, tail risk, adverse cases
D3	Admit a bounded specialist	Comparative success, security, recovery, named policy authority
D4	Renew, substitute or retire	Current baseline, drift, incidents, transfer tests and new proof

The output of D2 may be a failure receipt. The output of D3 is never an unlimited permission inferred from eloquence. The source definition of Specialist ASI requires superiority within a constituted mission and proof envelope, all hard gates, scoped authority, expiry and monitoring.[1] Trained models, a positive project NPV and a sophisticated mathematical appendix are not substitutes for those obligations. The retained estate should include the observations and rights map, supplied dependency graph, learned expressions and parameters, selected Beta, falsifiers, cost contracts, scenario witnesses, measurement calibration, code hashes and recovery instructions. Supplier substitution should be tested on the mission, not assumed from a file export. Preservation of memory must not silently preserve revoked operational permission.

22 Repairable scale and the physical price of continuity

The long horizon is a distributed scientific-industrial estate: energy collection, computation, instruments, material processing, fabrication and repair. The valuable quantity is service sustained through time, not nameplate capacity. An ideal sizing envelope is

$$P_{\text{useful}} = \eta A P_{\text{captured}} - P_{\text{auxiliary}}, \qquad N \geq \frac{P_{\text{useful}}}{A P_{\text{module}}}, \tag{22.1}$$

with efficiency η, service availability A and auxiliary consumption kept explicit. The retained screen sets auxiliary power to zero. At $P_{\text{useful}} = 10^{26}$ W, $\eta = 0.30$ and $A = 0.97$, it requires 3.436×10^{26} W captured, or 89.77% of the IAU nominal solar luminosity 3.828×10^{26} W.[22] This is captured incident power, not an existing generator. At 20% efficiency the same target exceeds one nominal star's budget. For illustrative 1-GW modules, this implies approximately 1.031×10^{17} modules. With mean up-time 3,650 days and mean repair duration ten days, an alternating-renewal account gives approximately 1.029×10^{16} repairs per year and 2.817×10^{14} simultaneous repairs. Adding a separate 2% common outage gives 97.732% mean availability under the retained assumptions.

Proposition 22.1 (Renewal imposes a throughput obligation). For N statistically stationary alternating-renewal units with finite mean up-time u and down-time d, long-run repair starts are $N/(u + d)$ per time unit and mean units down are $Nd/(u + d)$.

The renewal-reward identity proves both. Limited servers, correlated shocks and travel invalidate a literal independent-unit application; the new finite-source study quantifies one such departure. The stellar arithmetic is therefore a workload screen, not proof that enough repair stations, energy or materials exist. It must also include renewal of the repair infrastructure itself. If captured energy eventually thermalizes within the estate, it must be rejected as heat. A gray-radiator screen is $P_{\text{heat}} = \epsilon \sigma A_r (T^4 - T_s^4)$; a collector at nominal 1-AU flux requires $A_c = P_{\text{captured}}/1361$. These omit geometry, mutual irradiation, structural mass, parasitics, deterioration, transport and repair energy. They are necessary accounting questions, not sufficient designs. Energy exported outside the boundary must not also be charged as heat inside it.

23 Finite-speed authority and the mission portfolio

A remote controller needs at least a round-trip signal time $2D/c_0$. One astronomical unit gives approximately 998 seconds before computation, routing and actuation. A one-second local safeguard therefore cannot wait for that distant authority.

Figure 9. A newly generated conceptual horizon, cropped to exclude all generated claims and charts. Not an installation or engineering design.

Proposition 23.1 (A necessary causal control radius). If a feedback decision must complete within τ and local processing/actuation consumes $t_a < \tau$, then remote round-trip control requires

$$D \leq \frac{c_0}{2}(\tau - t_a). \tag{23.1}$$

This necessary bound does not account for noise, computation failure or controller adequacy.

The constructive response is a federation of local, qualified mission institutions with hard interlocks, preauthorized safe states, short authority leases and recoverable evidence. Long-horizon councils can reconcile goals, capital and obligations; they cannot overrule causality. Communication delay is not a reason for unconstrained local autonomy. It is a reason to specify the local envelope and failure behaviour before disconnection. The owner's next portfolio should therefore progress through measurable missions: conditional reserve availability; local repair readiness; correlated-defect detection; interchangeability and spares; robotic inspection; energy and thermal resilience; and only then new fabrication or remote assembly. Each successor may inherit code, observations and negative results, but must earn the new mission's proof. The potential scale is immense because these capabilities can support one another, not because any single solved equation authorizes all of them.

The continuity thesis

An institution becomes a durable productive asset when it can preserve useful service, reject its own inferior candidate, measure the event that matters and renew the physical capacity on which its next decision depends. The ambition is an expanding feasible world. The proof remains local, explicit and renewable.

24 From a research estate to a usable first mission

This Research & Use edition retains the original experiments, mathematical claims and negative comparisons. It adds a general field guide, a local mission-drafting workspace and a thin launcher for the unchanged v10.2 modules. These are usability and reproducibility contributions, not a new production release or a newly positive Alpha trial.

A useful first session

The owner leaves with one repeated decision, an accountable sponsor, observable outcomes, a fair incumbent/Beta comparison, permitted evidence, explicit exclusions and a named next reviewer. The engineer leaves with a minimal reproducible numerical example. Neither leaves with automatically granted authority.

Route	What the user does	Actual boundary
Owner	Complete and export the First Mission worksheet.	Draft completeness and supplied arithmetic; no verified evidence.
Practitioner	Inspect the frozen models and vary finite laboratory inputs.	Local synthetic calculations; no equipment connection.
Engineer	Run the native numerical example, then the neural-guided revision demo.	An unchanged scientific engine plus a new convenience adapter.
Verifier	Inspect commitments, tests, comparative records and missing qualification evidence.	Local review is not independent customer admission.

The tutorial exposes a dimensionless response $y = 1.8x + 0.25x^2$. Thirty-five training and fifteen selection observations form the development file. Twenty held-out and twenty stress observations are stored separately. Discovery receives only the development file; the challenge is passed to the native evaluator afterward. Coefficients and program structure are fitted through the real source interfaces. This is a deliberately simple file-contract demonstration, not a new physical law or protected test. The general workflow follows the supplied Constitution: constitute, reproduce baselines, form candidates, freeze, commission fresh proof, admit within a scoped envelope, monitor and renew. The diagrams on pages 7, 21 and 33 of that source respectively distinguish institutional decisions, the learning/serving split and authority levels. A draft schema is not the runtime policy schema. A template requires engineering and governance before it can control anything.[1]

25 The claim ledger is part of the user interface

A useful interface should not blend four different objects: an authored assumption, an executed local observation, an independently supported claim and an authorization. The new worksheet therefore reports only draft completeness and hypothetical net arithmetic. It exports no signature, no final proof receipt and no operational authority. Its amounts are user inputs, not calculated confidence bounds.

$$\hat{\alpha}_{\text{draft}} = \underbrace{12{,}000}_{\text{supplied lower spread}} - \underbrace{(3{,}000 + 4{,}000 + 2{,}000 + 1{,}000 + 1{,}000)}_{C+R+H+I+D} = 1{,}000 \text{ CAD.} \tag{25.1}$$

All amounts in this example are illustrative Canadian dollars per quarter. A required threshold of CA$5,000 is not met. Even a positive arithmetic threshold comparison would not establish the integrity of the lower bound, a fair baseline, rights, independent proof or hard-gate compliance. Costs already incorporated into utility must not be deducted a second time.

25.1 Reproduction is a route, not a deployment permit

The convenience launcher reports the installed environment without installing anything, routes a finite command set to native modules and preserves their no-overwrite behaviour. The guide uses an isolated Python environment and the chosen interpreter explicitly. Version pins identify the authoring reference; they do not constitute a complete transitive, hash-locked dependency assurance case. Python and pip document these distinctions.[23, 24] The field guide carries working templates for evidence, comparisons, experiments, proof handoff, weekly review, impairment and restoration. They are not a replacement for domain engineering or customer accountability. NIST's voluntary AI Risk Management Framework provides external risk-management context, not certification of this publication.[25]

A practical standard of excellence

The release is judged by faithful evidence, reproducible computations, usable instructions, explicit failure paths and honest unresolved limits. No self-awarded quality score establishes scientific truth. The accompanying release-review record states which checks were executed, which earlier results were preserved and which qualifications remain open.

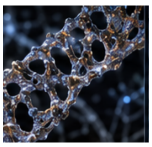

Figure 10. Conceptual image-generator material detail, editorial only; not a discovered material or measurement.

A Notation, dimensions and statistical units

Symbol	Meaning	Unit or boundary
a, l	Normalized age and load	Dimensionless; supplied observation domain
q_i, v_i	Interval unavailability / availability	Probabilities, not intensity rates
λ, μ	Failure / repair intensity	Per day; separately assumed in new repair study
A_j, D_j	Service fraction / demand	Probability and MWh per authored interval
Q, π	Markov generator / stationary law	Inverse days / probability vector
$v(S), x_i$	Coalition surplus / allocation	CAD per portfolio interval; transfers are not additional value
η	Scenario envelope or efficiency	Context-specific; not the same physical parameter
θ_i	Joint outage-and-reserve event	Probability per representative exposure
$\mathcal{F}, H$	Thermal exergy / physical storage	Joules when material parameters are calibrated
$D(p\|\pi)$	Relative entropy	Dimensionless nats; not heat by itself

One training cohort, one final portfolio and one observed exposure are different statistical units. The 624 annualization factor is a scale assumption. A finite design-catalogue optimum is not an optimum over every engineering configuration. Figures reporting mean service do not establish a bound on outage duration.

B Network and regret derivations

For $a_i, b_i \in [0,1]$,

$$\prod_{i=1}^{m} a_i - \prod_{i=1}^{m} b_i = \sum_{k=1}^{m} (a_k - b_k) \prod_{i<k} a_i \prod_{i>k} b_i, \tag{B.1}$$

so the absolute difference is at most $\sum_k |a_k - b_k|$. Apply this to the failure-product in a parallel block and to its shared series factors. An intervened component's dependence on the original forecast disappears in the source's action-effect formula, so retaining all ϵ_i is conservative. Multiply by nonnegative service-loss weights. For minimizers,

$$\begin{aligned} J(\hat{x}) - J(x^*) &= [J(\hat{x}) - \widehat{J}(\hat{x})] + [\widehat{J}(\hat{x}) - \widehat{J}(x^*)] \\ &\quad + [\widehat{J}(x^*) - J(x^*)] \le 2\Delta. \end{aligned} \tag{B.2}$$

If an approximate solver has declared optimization gap γ, add γ. If the feasible sets differ, the inequality is not applicable. Physical admission requires the true constraints to be satisfied, not a large financial penalty that merely makes violation uncommon in expectation.

C Repair generators, witnesses and limitations

The implementation forms Q, solves $Q^\top \pi = 0$ with one equation replaced by normalization, and separately compares the birth–death product formula with that linear solve. Product weights are evaluated in logarithms to avoid overflow. The tests check nonnegativity, normalization, generator row sums, stationary residuals, flow balance, rate-rescaling invariance, increased-crew stochastic ordering and the unlimited-server binomial limit. For the modulated model, each environmental block has its own failure and repair rates, and transitions between environments preserve the number failed. The stress marginal remains $0.002/(0.002 + 0.098) = 0.02$, while K becomes dependent on E. Service probability is not the product of this marginal with an unconditional repair distribution. This mechanism is precisely why a backlog can outlive the common outage. The new experiment enumerates 24 baseline staffing cases, 12 paid-hardening cases and 324 sensitivity designs. All outcomes are expectations under configured rates. The hardening effect is not learned or empirically verified. Finite-state numerical accuracy says nothing about unmodeled geographic travel, deferred maintenance, spare shortages or non-exponential durations.

D Core certificates and the meaning of a negative surplus

The least-core linear programmes use million-CAD scaling for conditioning and reconstruct all constraints in CAD afterward. Singleton constraints prevent negative allocations in a strict core, but their relaxation can make a diagnostic least-core allocation negative. The code reports both budget and coalition residuals. It never issues operating authority or transfers funds. If $v(N) < 0$ and $v(\{i\}) = 0$ for all members, the strict core is empty: nonnegative individual allocations cannot sum to a negative number. Even $v(N) > 0$ is insufficient; subgroup claims may jointly exceed the total. A designer can change pooling rules or permit subcoalition operation, but must then recompute the characteristic function. The 256 new LPs analyze the recorded values without modifying them.

E Joint evidence, stopping and economic monotonicity

For an intersection probability x, the four cells $(x, q - x, r - x, 1 - q - r + x)$ must be nonnegative. These inequalities produce the sharp Fréchet interval. If availability is observed only during normal operation, neither marginal r nor $r \mid O$ is identified without further sampling assumptions. The direct paired-event design records outage and reserve service at the same exposure time. The new synthetic campaign shares one common-outage sequence across all four sites. This creates dependence between simultaneous channels, which the Bonferroni argument permits. It does not permit serial dependence within each channel without adjustment. The observed event counts and fixed seed are recorded; repeated testing until a profitable bound appears would define a different, invalid fixed-sample protocol. Monotonicity follows from nonnegative coefficients $z_i M_i h_t/(1.08)^t$ multiplying each θ_i. Lower joint-service values therefore give lower NPV when every other coefficient is fixed. Taxes, prices, financing, risk aversion and future nonstationarity remain outside that statement. Confidence about four frequencies is not confidence about a whole business.

F Reproducibility map and fresh validation

Path	Content
`simulation/AEON_OMEGA_1_0/`	Complete retained customer simulation, models, code, observations, licences and first-party snapshot
`research/continuity_frontier.py`	New repair Markov chains, sensitivity study, entropy and robust least-core calculations
`research/joint_evidence.py`	Fréchet interpretation and new paired-event campaign
`research/results/`	Exact input assumptions, distributions, allocations and decisions
`latex/`, `figures/`	Bilingual manuscript, bibliography and numerical vector figures
`assurance/`	Fresh tests, replay and delivery checks

A fresh same-host run of the retained study reproduced all 134 primary files and all four final secondary files byte-for-byte. The 55 retained extension tests and 52 new mathematical tests passed. The new calculations are also replayed into a fresh directory; detailed results and delivery checks are recorded in the machine-readable validation report. These checks are not an independent audit. The inherited full-product 840-test regression is not claimed as newly rerun in this publication.

```
python3 VERIFY_PUBLICATION.py
python3 research/continuity_frontier.py --out var/new-frontier
python3 research/joint_evidence.py --out var/new-joint.json
python3 -m pytest research/tests -q
python3 simulation/AEON_OMEGA_1_0/REPRODUCE.py --out var/original-replay
```

Use a fresh compatible Python 3.13 environment and the pinned research dependencies; no external API is required. Rebuilding the papers requires XeLaTeX, Biber and locally installed fonts documented in the build guide. Font files are not distributed. Exact bitwise replay was checked on the authoring host, not across every numerical backend.

G Limitations and a falsifiable research agenda

The largest unresolved questions are measurement validity, causal intervention effects, unknown dependencies, nonstationarity and feasible field recovery. The small numerical models are chosen for inspection, not realism by assertion. No neural model discovers the dependency graph or the hardening effect. No new material, universal theory, orbital installation or independently qualified customer successor is produced. A stronger next study should collect paired outage/reserve evidence and timed repair histories; challenge the learned intervention effects; include travel, spares and heterogeneous units; compare available suppliers at real costs; prospectively freeze both model and statistical protocol; and test a restore-and-substitute drill. The research should publish losing outcomes, unserved service and no-build decisions alongside successful ones. The durable promise is not that the current candidate is the ceiling. It is that a customer can retain a disciplined means of replacing the wrong answer, qualifying a better one, and keeping the physical estate useful while that succession occurs.

H Complete primary cohort ledger

The table retains cohort-level annualized net differences rather than treating every portfolio as a separately trained replication. All values subtract the same CA$9 million incremental annual cost.

Seed	Native fits	Generic Beta, M CAD	Informed Beta, M CAD
9100	3 243	-9.000	-9.000
9101	3 137	-9.000	-9.000
9102	3 182	-9.007	-9.007
9103	3 141	-5.980	-9.000
9104	3 141	-9.000	-9.000
9105	3 184	-8.953	-9.000
9106	3 138	-8.996	-9.000
9107	3 104	-9.000	-9.000

A negative number is not erased by a strong explanation or a larger incumbent contrast. Reusing the evidence to choose a different component can still be a valuable institutional outcome.

References

[1] V. Boucher. *Neural-Symbolic SUCCESSOR Ω: Customer-Owned Mission Intelligence*. Supplied corporate white paper. Architectural definition; not empirical certification. MONTREAL.AI, Aug. 25, 2026.

[2] MONTREAL.AI. *NEURAL-SYMBOLIC SUCCESSOR Ω, version 10.2.0*. Supplied implementation. Original archive SHA-256 begins fb48474666ca39b6; 55 unchanged first-party files retained in AEON. Sept. 5, 2026.

[3] V. Boucher. *AEON Ω: The Continuity Estate, research simulation 1.0*. Supplied executable research: eight synthetic cohorts, portable artifacts and conditional economics. Included in simulation/AEON_OMEGA_1_0/. MONTREAL.AI and QUEBEC.AI, Sept. 5, 2026.

[4] 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. URL: https://arxiv.org/abs/2402.12275.

[5] NIST/SEMATECH. *e-Handbook of Statistical Methods: Parallel or Redundant Model*. URL: https://www.itl.nist.gov/div898/handbook/apr/section1/apr183.htm (visited on 09/05/2026).

[6] NASA. *Prognostics Center of Excellence Data Set Repository*. Future data context only; no NASA data used in the retained experiment. URL: https://www.nasa.gov/intelligent-systems-division/discovery-and-systems-health/pcoe/pcoe-data-set-repository/ (visited on 09/05/2026).

[7] N. K. Jaiswal. *Finite-Source Queuing Models*. Technical Memorandum 45. Case Western Reserve University, 1966. URL: https://commons.case.edu/wsom-ops-reports/210/.

[8] D. Gross and J. F. Ince. « The Machine Repair Problem with Heterogeneous Populations ». In: *Operations Research* 29.3 (1981), pp. 532–549. DOI: 10.1287/opre.29.3.532.

[9] 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.

[10] U. Seifert. « Stochastic Thermodynamics, Fluctuation Theorems, and Molecular Machines ». In: *Reports on Progress in Physics* 75 (2012), p. 126001. DOI: 10.1088/0034-4885/75/12/126001. arXiv: 1205.4176.

[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, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada. « Control Barrier Functions: Theory and Applications ». In: *European Control Conference*. 2019, pp. 3420–3431. DOI: 10.23919/ECC.2019.8796030. arXiv: 1903.11199.

[14] L. S. Shapley. *A Value for N-Person Games*. P-295. RAND Corporation, 1952. URL: https://www.rand.org/pubs/papers/P295.html.

[15] H. E. Scarf. « The Core of an N Person Game ». In: *Econometrica* 35.1 (1967), pp. 50–69. DOI: 10.2307/1909383.

[16] M. Maschler, B. Peleg, and L. S. Shapley. « Geometric Properties of the Kernel, Nucleolus, and Related Solution Concepts ». In: *Mathematics of Operations Research* 4.4 (1979), pp. 303–338. DOI: 10.1287/moor.4.4.303.

[17] 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.

[18] 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.

[19] 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.

[20] C. J. Clopper and E. S. Pearson. « The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial ». In: *Biometrika* 26.4 (1934), pp. 404–413. DOI: 10.1093/biomet/26.4.404.

[21] 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.

[22] A. Prša, 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] Python Software Foundation. *venv: Creation of virtual environments*. Python 3.13 documentation; setup context only. URL: https://docs.python.org/3.13/library/venv.html (visited on 09/06/2026).

[24] pip developers. *Repeatable Installs*. URL: https://pip.pypa.io/en/stable/topics/repeatable-installs/ (visited on 09/06/2026).

[25] National Institute of Standards and Technology. *AI Risk Management Framework*. Voluntary risk-management context; no certification implied. URL: https://www.nist.gov/itl/ai-risk-management-framework (visited on 09/06/2026).