MONTREAL.AI QUEBEC.AI

SUCCESSOR Ω — ANABASIS Ω

SCIENCE · OPTIONS · EVIDENCE · CAPACITY

The Expansion Institution

Buy the evidence. Govern the options.
Build the next capability.

Vincent Boucher President, MONTREAL.AI & QUEBEC.AI

ANABASIS-TR-01 / 1.4 6 SEPTEMBER 2026

AI-GENERATED CONCEPT ART · NOT AN ENGINEERED DESIGN

Publication record and evidence contract

Title	SUCCESSOR Ω — ANABASIS Ω
Subtitle	The Expansion Institution: buying evidence, governing options, building the next capability
Author	Vincent Boucher, MONTREAL.AI & QUEBEC.AI
Edition	ANABASIS-TR-01 / Illustrated Research Publication 1.4 / 6 September 2026
Empirical basis	Retained ANABASIS simulation 1.0; four authored synthetic cohorts
Implementation	Unchanged SUCCESSOR Ω v10.2 scientific engine; customer-side adapters
Current-use reference	SUCCESSOR Ω v11.0.0 is a separate release, not a relabelling of the retained trial
Status	AI-assisted research manuscript; no independent peer review or production admission

The institutional proposition

Own the capacity to determine what is worth proving next, preserve what reality refused, and convert accepted outcomes into the next useful capability.

The paper joins a reproducible numerical study to a larger scientific and institutional argument. It preserves unsuccessful candidates and the absence of proprietary advantage over a serious alternative. New publication-side calculations concern evidence access, conditional risk and physical renewal. Their assumptions are disclosed separately; they do not retroactively improve the earlier experiment.

Executed means local computation on synthetic observations. **Derived** means a statement proved under explicit assumptions. **Conditional** means an authored physical, commercial or statistical scenario. **Prospective** means an unexecuted real-world requirement. These classes are not interchangeable.

The architectural sources define Specialist ASI as an earned mission-bounded status and SUCCESSOR Ω as a complete admitted institution, not merely a fitted model. This study supplies research components and local records, not independent proof, an authority envelope or customer cash.[1–3]

The visual sequence uses five distinct AI-generated scenes and two source-authored vector schematics; no scene is reused within an individual-language paper. These are conceptual illustrations, not photographs, apparatus specifications or empirical evidence. Corresponding English and French openings intentionally share the same imagery. Exact typeset captions, equations and tables control the claims. Original image microtext has no scientific authority. Edition 1.4 changes editorial presentation, not the retained research results.

Original software and third-party licences remain applicable. Exported files do not by themselves establish legal title, independent custody or permission to use restricted observations. No signature, external verifier, deployment, financial receipt or universal discovery capability is fabricated.

Abstract

A wealthy owner can finance a productive asset before knowing whether it deserves to exist. SUCCESSOR Ω — ANABASIS Ω studies the institution needed to purchase decision-changing evidence before irreversible commitment. The retained synthetic study combines an unchanged neural-symbolic scientific engine with a finite partially observed Mission Gym, a material-and-cash ledger and a customer-authored sequential planner. Four cohorts train four neural guides and twelve symbolic generations through 8,935 native candidate fits. Neural ablation changes every proposed calibration probe. All four second-generation models pass the nominal local protocol, all fail a changed process, and three third-generation candidates recover local validity.

In the authored twelve-year project, the contingent policy buys a low-cost recovery coupon before deciding whether a calibration blank or pilot is worth purchasing. Expected project net present value after assays is CA$104.664 million, compared with CA$53.872 million for one-step research and CA$96.751 million for information per dollar. It declines construction in 47.86% of outcomes. An equally informed conventional predictor using the same complete planner ties the physical decisions; an additional CA$12 million programme allowance therefore creates no proprietary premium. An independence-only observation model reaches inadmissible investment branches with 58.08% probability despite a higher mean value. A further adverse change to the pilot defeats the nominal safety claim.[4]

Pruning investment leaves against four pilot-contamination models reduces nominal expected value to CA$78.345 million. Edition 1.3 separately reoptimizes nominal value under a persistent-channel risk family: value rises to CA$86.901 million and expected assay cost falls to CA$1.054 million. An endpoint argument establishes the stated interval coverage; this is not worst-case expected-value optimization. A six-player evidence-access game has a core, but its Shapley allocation violates twelve coalition constraints. An exact counterexample separates the value of information under fixed actions from mandatory posterior-dependent admissibility.

The formal analysis connects finite belief-state planning, decision-sensitive uncertainty, Bayesian information, proper scoring, audit incentives, statistical claim budgets, material balance, free energy, feedback thermodynamics and age-structured renewal. A favorable, separate capital route returns contributed capital before funding science and replacement. The stellar horizon remains a physical research programme with power, heat, feedstock, retirement and control latency explicit. The contribution is an inspectable path from learned hypotheses to governed options and paid capacity, not unrestricted theory discovery, independent Specialist ASI, realized customer Alpha or production authority.

KEYWORDS Neural-symbolic models; sequential experimental design; shared calibration; Bayesian games; Mission Gym; proof capital; material recovery; renewal; institutional succession.

Owner decision brief

The best useful action is not always to build

Constitute one consequential decision. Buy the smallest evidence package that can change it. Preserve the strongest suitable alternative. Release capital only through the appropriate review, not because a simulation reports a large number.

Question	Retained or new result
Contingent project value after assay costs	CA$104.664 M
Improvement over one-step research, within that value	CA$50.791 M
Chance of no construction	47.86%
Worst nominal reached-leaf severe-loss probability	7.72%
Programme comparison with informed OLS and full planner	CA$−12.000 M
New finite-model guard: nominal project value	CA$78.345 M
Edition 1.3 replan: nominal value / expected assay cost	CA$86.901 / 1.054 M
New information game: Shapley coalition violations	12 of 62
Real customer Alpha credited / production authority	0 / none

The first five rows belong to one authored project and its distinct comparisons. The guard, game and Edition 1.3 replan are separately scoped publication-side calculations. No amounts should be added across these columns or to prior AEON, DAEDALON or ASTERION studies.

Keep the learned programs, benchmark, latent-bias model, failed successors, contingent policies and source-linked records. **Do not promote** an exclusive-model premium, pilot independence, a real sampling guarantee or an admitted Specialist ASI. **Investigate next** calibration quality, adverse observation channels, a procurable baseline and the inspection bottleneck before adding feed capacity.

Three routes through the publication

Owner: Sections 1, 6–9, 17–23 and the practical guide; Appendix I gives the new decision. **Scientist:** Sections 3–5 and 10–16, followed by the derivations. **Engineer and verifier:** the source map, retained experiment, new tests, exact commands and qualification register. The bundled general Field Guide supplies the wider operating method.

The numerical adapter is a research tool. A mission may usefully end in rejection, uncertainty, a better baseline or a smaller experiment. A label, image, positive expected value or signed export cannot replace independent evidence and accountable admission.

Contents

Publication record and evidence contract i

Abstract ii

Owner decision brief iii

I / THE OWNER AND THE MISSION — The next capability 1

1 The expansion mandate 2

2 From SEIZE to a complete institution 2

3 A learned world program that changes a decision 3

II / EVIDENCE AND DECISIONS — The price of knowing 4

4 A hidden-state mission world 5

5 Material, energy and value must reconcile 5

6 The finite contingent decision problem 6

7 The contingent policy and its value 6

8 Correlated evidence changes the permitted investment 7

9 What a finite-model safeguard costs 8

III / INFORMATION AND PHYSICAL REALITY — The limits that create value 10

10 Information is not the same as decision value 11

11 Decision-sensitive model error 11

12 Bayesian geometry without metaphysical claims 12

13 Recovery physics: conservation, separation and exergy 12

14 Physical feedback and institutional free energy 13

IV / STRATEGY, PROOF AND OWNERSHIP — The institution that can say no 15

15 The game of access to evidence 16

16 Incentives for useful measurements, not favourable verdicts 17

17 Mission Alpha and successive opportunities to be wrong 17

V / CAPITAL AND THE PRODUCTIVE HORIZON — Build the next feasible world 19

18 Capital comes back before surplus compounds 20

19 Capacity follows the binding constraint 20

20 A growing estate must also replace itself 21

21 A stellar-scale horizon with a physical denominator 21

VI / THE OWNER'S OPERATING METHOD — From a first mission to a successor 23

22 Using SUCCESSOR Ω generally 24

23 Version-specific use and truthful portability 24

24 The complete handoff and admission boundary 24

25 Limitations and the expansion thesis 25

A Notation and measurement boundaries 26

B Exact terminal account and loss definition 26

C Posterior sufficiency, ordering and scope 27

D Finite robust guard: complete interpretation 27

E Coalition certificate and its proper use 27

F How the source documents are used 28

G Reproduction and provenance 28

H What to do next 29

I From pruning decisions to redesigning research 30

The next capability

Intelligence becomes an enduring asset when it helps determine what deserves to be built next.

CONCEPT ART OR SCHEMATIC · NOT AN INSTALLATION OR EVIDENCE

1 The expansion mandate

Dr. Soraya Delcourt is a fictional computational physicist directing a CA$620 billion estate with CA$18 billion of group liquidity. Neither amount belongs to the real user. The setting permits a consequential research question without implying that every investment is affordable, profitable or authorized. Her ring-fenced mandate allows at most CA$750 million of construction, CA$6.5 million of assays and three distinct evidence actions.

The choice is recovery, fabrication, an integrated line, or no new project. Construction is paid at time zero. The asset earns nothing for two commissioning years. A twelve-year pre-tax NPV uses an authored 8% discount rate. The planner must decide not just what to build but what to learn before choosing.[4]

The owner's instruction

Find the next capability worth proving. Buy the smallest useful observation. Count only accepted material and external sales. Preserve every reason to stop.

AEON addressed continuity of an existing estate. ANABASIS moves to capability formation: which productive addition warrants investigation and capital? The earlier monetary results are not inputs to this study. A new mission requires a new objective, physical account, comparator and proof envelope.

Three assets must be separated. A *predictor* estimates a response; a *decision system* chooses among permitted actions; an *institution* supplies rights, independent challenge, accountable ownership, recovery and renewal. The owner may keep the latter two while replacing the first. That substitution is a strength, not a concession, when a conventional component serves the mission equally well.

2 From SEIZE to a complete institution

The controlling progression is Successor Manifold, Mission Advantage Gradient, SEIZE Underwriting, Bounded AGI Jobs, Specialist ASI, then accountable Mission-Sovereign admission. In this study it is an architecture-to-artifact map, not a claim that all institutional stages ran.[1, 2]

Surface the changed opportunity and credible alternatives. **Evaluate** the value of learning before commitment. **Instantiate** a bounded research option. **Zero in** on one frozen candidate and the evidence it needs. **Elevate** a review request. Elevation is not a release certificate.

The five decision functions—underwrite, execute, accept, admit and allocate—must not be confused with corporate phase gates or action authority. In particular, simulation permission does not authorize procurement, apparatus control or investment. In the current package, actor fields and a local hash chain are useful records but not an adversarial security boundary.

Layer	Executed here	Remaining obligation
Native science	Bounded search, local neural training, executable expressions	Customer measurement validity
Customer Gym	Hidden states, observations, costs, stop actions	Protected custody and field fidelity
Decision system	Finite contingent planning and ledgers	Procurable comparator, actual costs
Institution	Source map, qualification gaps, failure records	Independent acceptance and admission

The four-phase state machine on PDF page 72 of the Gym source makes this separation explicit. Its arrows require an actor, exact evidence, relevant versions, an authority record and rollback route. A diagram is not evidence that an arrow occurred. The 21-job map in the delivered study therefore leaves independent proof and operational admission open.

3 A learned world program that changes a decision

The native engine receives normalized intensity $u \in [0, 2]$, contamination $z \in [0, 1]$ and a dimensionless recovered-mass fraction. Initial data cover a narrow range. Later data broaden it, with uniform synthetic measurement noise of amplitude 0.0002. The engine searches a finite expression language with explicit budgets; it does not derive a chemical process from first principles.

Sparse equation discovery and neural-assisted symbolic regression provide relevant methodological precedents; executable program-based environment models supply another route. These papers establish research mechanisms, not integrated SUCCESSOR Ω performance.[5–7]

Each of four cohorts has 32 initial training and 20 initial selection observations. The widened nominal formation has 101 training points, including the new probe, and 48 selection points; 64 holdout and 48 stress points are separate. A shifted process has 120 new training, 48 selection, 64 holdout and 48 stress observations. Partition separation is local and author-controlled, not independent proof custody.

Four 12-hidden-unit neural guides are trained for 1,000 steps each. Their disagreement with the current symbolic hypothesis influences the selected probe. An ablation changes that choice in all four cohorts; this establishes influence, not globally optimal experimentation. Across three generations, 8,935 candidate fits produce twelve retained symbolic artifacts.

The first G2 approximation is

$$\hat{y}(u, z) = 0.401267 - 0.229983z + 0.378340u - 0.050054zu + 0.136798 \cos u - 0.038130e^{u}. \tag{1}$$

The authored nominal response is quadratic. Multiple expressions approximate it over the permitted region; the fitted trigonometric and exponential terms are not a newly discovered causal law. All G2 candidates pass their nominal protocol, all fail a changed process, and three G3 candidates recover local validity. The fourth failure remains in the archive. No successor inherits authority.

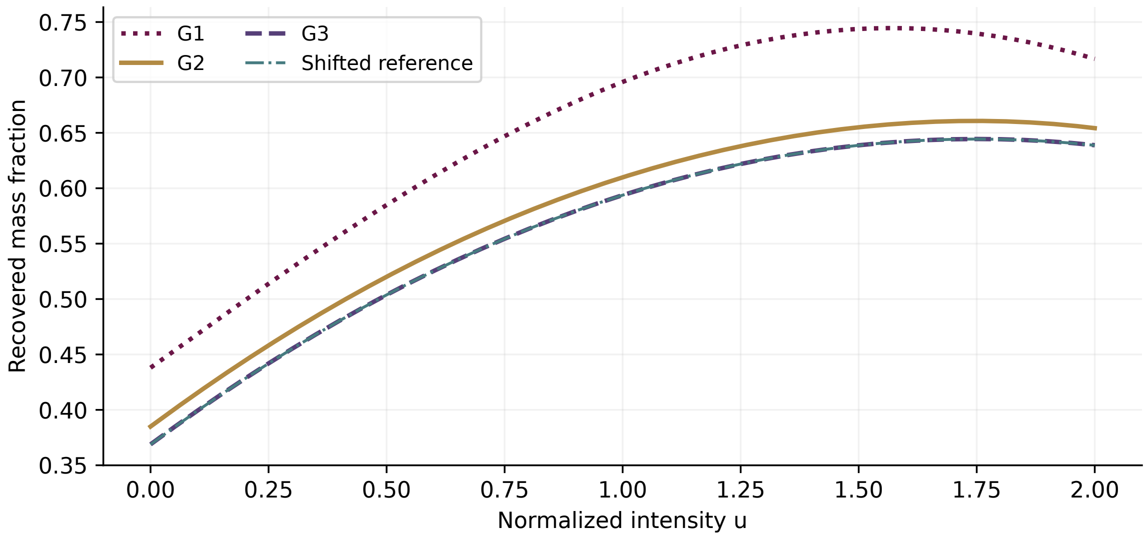

Figure 1 – *Retained cohort 64100 at contamination 0.5. Programs are evaluated directly from their stored artifacts; curves describe an authored response, not measured industrial yield.*

The price of knowing

A convincing observation is valuable only through the decision it can legitimately change.

CONCEPT ART OR SCHEMATIC · NOT AN INSTALLATION OR EVIDENCE

4 A hidden-state mission world

Let $\theta = (r, f, b) \in \{0, 1\}^3$ denote recovery suitability, fabrication suitability and persistent assay bias. The prior factors into probabilities 0.70, 0.30 and 0.35. Conditional on this full state, observations are independent; marginally, repeated coupons share the latent bias and are dependent.

A history $h = \{(a_k, o_k)\}$ contains test identities and binary outcomes. The posterior is

$$b_h(\theta) = \frac{b_0(\theta) \prod_k L_{a_k}(o_k \mid \theta)}{\sum_\vartheta b_0(\vartheta) \prod_k L_{a_k}(o_k \mid \vartheta)}. \tag{2}$$

The user supplies this observation structure. The learned yield function does not autonomously discover the calibration state or prove the assays' physical independence.

Test	Cost, CA$ M	Nominal positive probability
Recovery A	0.65	0.88 if biased; otherwise 0.82 suitable, 0.14 unsuitable
Fabrication	0.65	Same pattern for fabrication suitability
Recovery B	0.45	Same channel as Recovery A, not new independent calibration
Joint pilot	4.50	0.96 if both suitable; 0.12 if exactly one; 0.025 if neither
Calibration blank	0.35	0.95 if biased, 0.045 otherwise

This is a small partially observed control problem in the sense of belief-state planning.[8] Finite observation histories are enumerated exactly. A local Python field hides state from the acting policy, but a claimant with source access can inspect the simulator; such hiding does not constitute independent protected evaluation.

5 Material, energy and value must reconcile

The twelve terminal choices comprise no-build; recovery and integrated configurations at five intensities; and fabrication alone. Recovery processes 180,000 tonnes annually. Its contamination is 0.10 in a suitable state and 0.95 otherwise. The authored response is

$$y(u, z) = 0.50 + 0.34u - 0.09u^2 - 0.23z - 0.05uz. \tag{3}$$

For feed F, recovered mass is Fy and residue $F(1 - y)$. An integrated line internally transfers $T = \min\{110000, Fy\}$ tonnes to fabrication. Fabrication conversion is 0.92 or 0.51, conditional on suitability. Recovered material transferred internally cannot also count as external sales.

Proposition 1 (Accounted mass and external sales)**.** With imported fabrication feed I, finished output $G = \eta_f(I + T)$, fabrication scrap $S = (1 - \eta_f)(I + T)$ and externally sold recovered mass $R_s = Fy - T$,

$$F + I = F(1 - y) + S + R_s + G. \tag{4}$$

External sales are $p_R R_s + p_G G$, without a revenue term for T.

Proof. Substitute the four definitions. The internal transfer cancels. Counting it as another sale would record the same mass on both sides of the commercial boundary. □

Energy is $F(0.30 + 0.45u^2) + 1.3(I + T)$ MWh in the authored process account. Prices, fixed cost, disposal and commissioning are declared inputs, not discovered constants. The checker reconstructs the arithmetic separately from the planner but shares the authored physical model. It detects 480 injected ledger mutations across 96 state/design accounts; it cannot validate an incorrect model of nature by agreeing with itself.

6 The finite contingent decision problem

Let $U_d(\theta)$ be a terminal project's NPV before assays. Its capex is K_d. A terminal decision is admissible when $K_d \le 750$ million and

$$b_h^\top \mathbf{1}\{U_d < -100\} \le 0.10, \tag{5}$$

with monetary quantities here in millions of Canadian dollars. No-build has zero project value and is always available. Assay costs are subtracted from utility, but the retained loss event deliberately excludes them and the separate programme allowance. A different loss definition is a different constitution.

Let $\mathcal{T}(h)$ contain distinct affordable tests remaining before the three-test limit. Then

$$\begin{aligned} V(h) = \max\Bigg\{ & \max_{d \in \mathcal{D}(h)} b_h^\top U_d, \\ & \max_{a \in \mathcal{T}(h)} \left[-c_a + \sum_{o \in \{0,1\}} P(o \mid h, a) V(h \cup \{(a, o)\}) \right] \Bigg\}. \end{aligned} \tag{6}$$

Proposition 2 (Exactness inside a finite mission)**.** For a fixed observation model, payoff table, admissibility rule and finite test budget, backward induction in (6) gives an optimal deterministic policy among the permitted history-dependent policies.

Proof. At terminal histories, maximize among permitted terminal actions. Assume successor histories are solved. Taking each affordable test yields its known cost plus the conditional expectation of already optimal continuation values. The maximum over those choices and stopping is optimal. Induct backward to the root. This proves neither the model nor the adequacy of the finite catalogue. □

The representative search visits 131 cached history states. Budgets, unavailable tests and no-build are part of the optimization, not penalties to be waived after seeing the result. The retained information-per-cost policy is an explicit heuristic that can keep testing to its limit; it is not an exhaustive characterization of all information-based methods.

7 The contingent policy and its value

The first action is recovery coupon B, costing CA$450,000. A positive B leads to the calibration blank. A negative blank supports a recovery recommendation; a positive blank requires the pilot. A negative B leads to A: a second negative stops, while a positive A requires the pilot. Pilot outcomes separate the remaining investment and no-build branches.

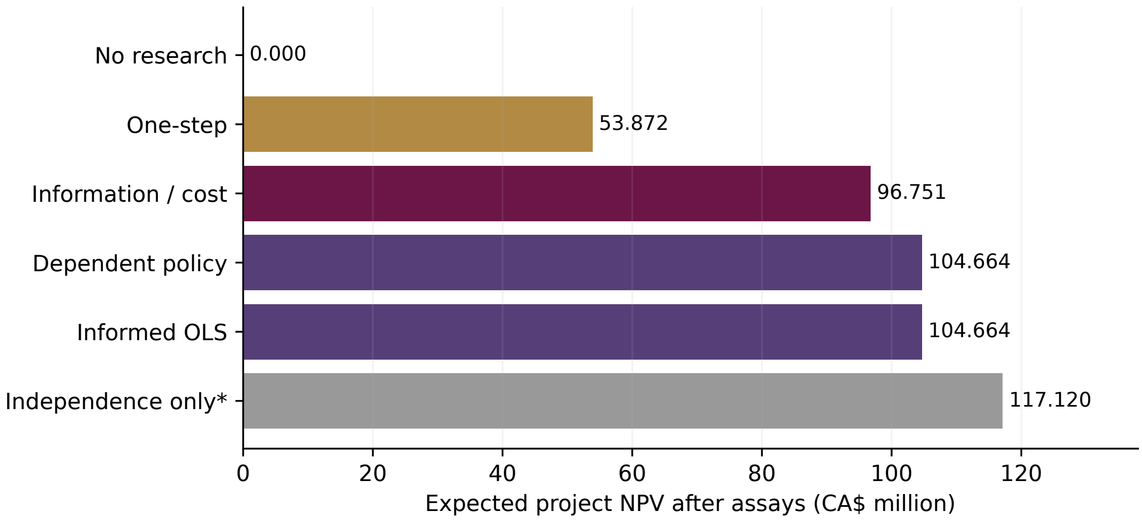

Figure 2 – *Exact expected NPV within the authored project after assay costs. The high independence-only value is inadmissible under the nominal loss constraint. The informed OLS comparator ties the complete method.*

Policy	NPV, CA$ M	Expected assay cost, CA$ M
No research / no-build	0.000	0.000
One-step value	53.872	4.500
Information per dollar	96.751	2.676
Dependent contingent policy	104.664	2.859
Informed OLS + same planner	104.664	2.859

The candidate declines construction in 47.86% of outcomes. Its maximum nominal reached-leaf severe-loss probability is 7.72%, not zero. The expected loss-event probability across all outcomes is 3.73%. These are exact probabilities of an authored eight-state world, not confidence estimates for future facilities.

The CA$50.791 million difference from one-step research is contained within CA$104.664 million. It is not additional wealth to add to it. Four different fitted candidates choose the same central policy; this is useful stability within the selected scenario, not a representative industrial population study. The 27 retained sensitivity settings and seven exploratory configurations are provided without transforming scenario coverage into a statistical confidence claim.

8 Correlated evidence changes the permitted investment

If the common calibration state is integrated out separately for each observation, the planner treats repeated coupons as independent witnesses. After two positive recovery assays, the unsuitable-recovery probability becomes 8.80%. Retaining the shared state yields 14.66% instead. The calibration-bias posterior rises to 46.67%; the independence-only model never learns it and leaves it at 35%.

Proposition 3 (Marginal dependence from a shared channel)**.** If two identical binary tests are independent given binary bias B and suitability R, then conditional on R,

$$\mathrm{Cov}(O_1, O_2 \mid R) = p_B(1 - p_B)(\ell_1 - \ell_0)^2 \geq 0, \tag{7}$$

where $\ell_b = P(O_i = 1 \mid B = b, R)$ and $p_B = P(B = 1 \mid R)$.

Proof. Apply total covariance. Conditional covariance given B, R is zero; covariance of conditional means is the variance of a two-valued variable. □

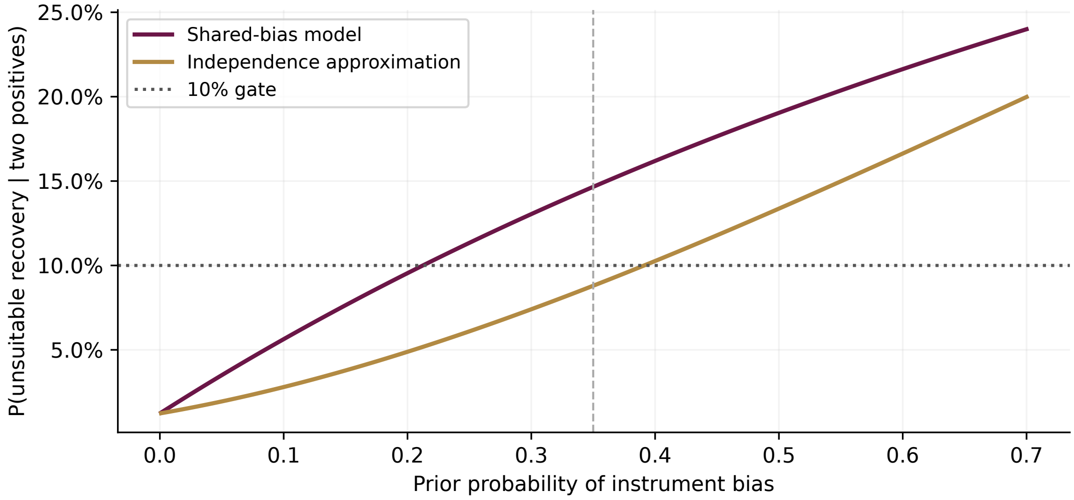

Figure 3 – *A shared calibration state can move the same pair of positive observations across the frozen 10% risk gate. Computed posteriors are conditional on the stated channels.*

The independence-only policy obtains CA$117.120 million expected value in the nominal reference but reaches leaves violating (5) with 58.08% probability. A higher objective is not a waiver of a hard constraint. This is Gym basis risk, not a failure repaired by a more flattering chart.

9 What a finite-model safeguard costs

The retained policy also fails when the nominally separate pilot is contaminated by bias: one reached investment leaf has severe-loss probability 29.27%. The new publication calculation freezes the original test sequence and examines four specified biased-pilot positive probabilities: 0.12, 0.30, 0.60 and 0.84. At an investment leaf, it switches to no-build if any member of this finite set breaches the loss gate.

Proposition 4 (Scoped leaf guard)**.** With identical terminal loss events and correct posteriors for each member of a stated finite model set, replacing every violating investment leaf by zero-loss no-build makes all retained investment leaves satisfy the gate for every member of that set.

Proof. A retained investment passes each model's test by construction. Every replaced leaf has no-build loss probability zero. The test does not assert anything about a model outside the stated set. □

Nominal expected NPV falls to CA$78.345 million, a CA$26.318 million reduction. No-build probability rises to 60.22%. Expected assay expenditure remains CA$2.859 million because the guard leaves test choices unchanged. The procedure is conservative, not an optimized robust policy: some tests may become wasteful once their investment branch is rejected.

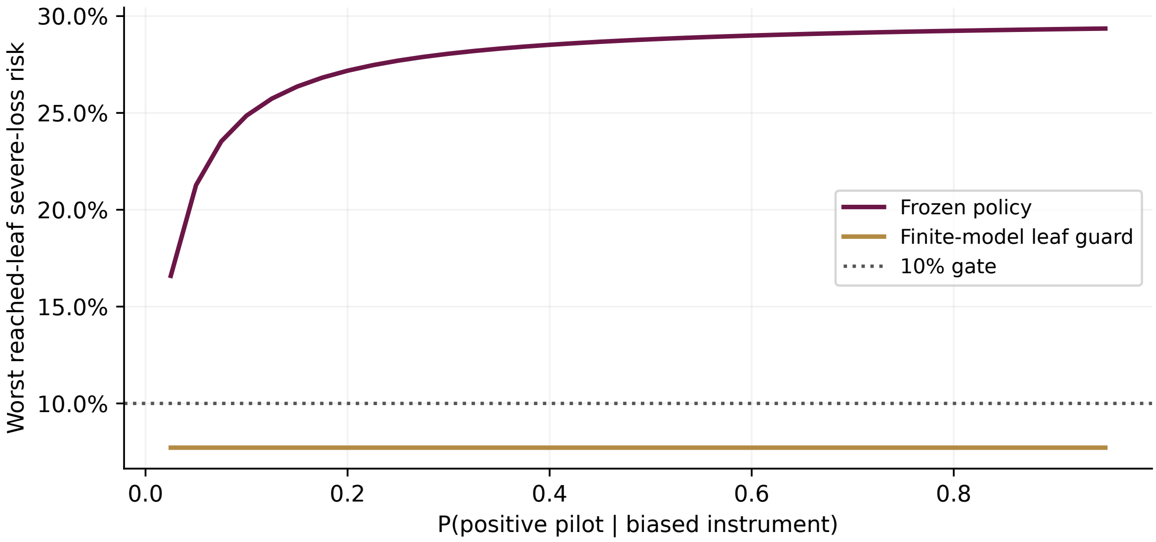

Figure 4 – *A 38-point pilot-channel sensitivity of the frozen and guarded policies. Only the four declared anchor models define the guard's stated scope; the grid is a diagnostic, not a probability distribution or confidence region.*

Robust dynamic programming requires care about the ambiguity set. Allowing an adversary to select a new worst conditional law at every history is a different problem from one persistent but unknown instrument regime. Rectangularity assumptions cannot be omitted when a minimax Bellman recursion is invoked.[9]

The limits that create value

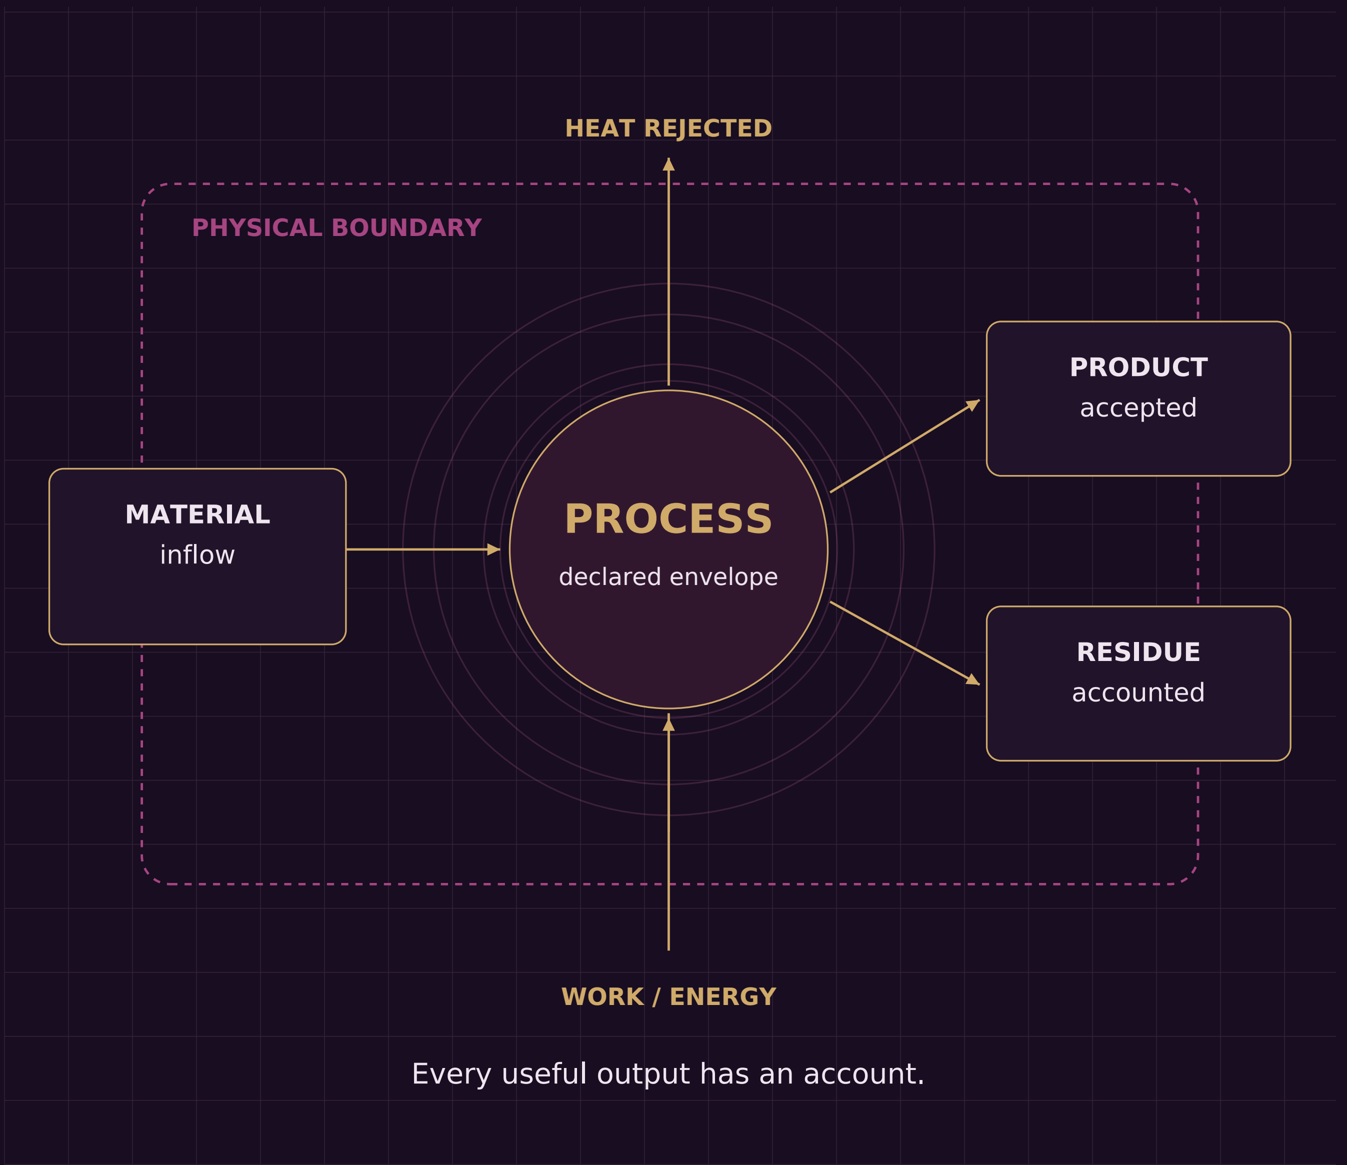

Information changes feasible decisions. It does not abolish the costs of measurement, work, matter or time.

10 Information is not the same as decision value

For a fixed terminal action set, let $W(b) = \max_d b^\top U_d$. Because a maximum of linear functions is convex, a cost-free observation with posterior martingale $\mathbb{E}[b_Z] = b$ has nonnegative expected value of information.

Proposition 5 (The fixed-action information condition)**.** If the decision maker may use the same actions after observing Z, then $\mathbb{E}[W(b_Z)] - W(b) \geq 0$. An observation that never changes a maximizing action can have positive mutual information and zero decision value.

Proof. Jensen's inequality gives the first claim. If one action d^* maximizes for every reachable posterior, $\mathbb{E}[W(b_Z)] = \mathbb{E}[b_Z]^\top U_{d^*} = b^\top U_{d^*}$. □

This is the relevant restriction behind comparison-of-experiments results.[10] It is not valid to quote information monotonicity while silently changing the permitted action set after each signal.

Proposition 6 (A mandatory posterior gate can break that monotonicity)**.** There exists a zero-cost, Bayes-plausible signal that reduces the value of a decision problem with mandatory full-posterior chance constraints.

Proof. Investment pays 100 in a good state and -110 in a bad state; no-build pays zero. The prior bad probability is 0.09, with admissibility requiring at most 0.10 probability of a payoff below -100. Investment is initially permitted and worth 81.1. A signal yields posterior bad probabilities 0.11 and 0.01 with probabilities 0.8 and 0.2. Their mean is 0.09, so the experiment is Bayes-plausible. The first branch prohibits investment. The second permits value 97.9. Expected constrained value becomes $0.2(97.9) = 19.58$. The decrease is 61.52 despite zero observation cost. The explicit signal channel is supplied in the new calculation. □

This is not an argument for suppressing safety-relevant evidence. It shows that research, disclosure and action-admissibility rules must be constituted together. Optional tests remain optional before purchase; once a material observation is known, the institution cannot pretend not to know it merely to preserve an attractive recommendation.

11 Decision-sensitive model error

Prediction error matters through payoffs and feasible choices. Let the true payoff vector for a fixed terminal decision be U_d, with approximation $\widehat{U}_d$ satisfying $\|U_d - \widehat{U}_d\|_\infty \leq \varepsilon$. Under a fixed policy, averaging over states and histories cannot amplify that terminal payoff bound. But small error can change the threshold event $\{U_d < -100\}$ discontinuously.

Proposition 7 (Regret bound with fixed admissibility)**.** If $|J(\pi) - \hat{J}(\pi)| \leq \varepsilon$ for every policy in the same admissible class, and $\widehat{\pi}$ is δ-optimal for $\hat{J}$, then

$$J(\pi^*) - J(\widehat{\pi}) \leq 2\varepsilon + \delta. \tag{8}$$

Proof. Add and subtract $\hat{J}(\pi^*)$ and $\hat{J}(\widehat{\pi})$. The two approximation terms are bounded by ε and the optimization term by δ. □

To transfer a loss gate, a robust sufficient condition is $P(\widehat{U}_d < -100+\varepsilon) \leq 0.10$, since $U_d < -100$ implies $\widehat{U}_d < -100+\varepsilon$. An average test error does not provide the uniform bound required here. A discontinuous action change near a capital, budget or risk boundary must therefore be evaluated at the decision level.

The same distinction applies to learned yield. On a purely recovery configuration, annual external sales vary linearly with recovered tonnes, but integrated capacity saturation introduces a kink. A regression fit can be accurate on average while selecting the wrong side of that kink. The finite action catalogue and separate ledger make such decision consequences inspectable; they do not remove model risk.

12 Bayesian geometry without metaphysical claims

Updating the belief state changes a distribution over mission hypotheses, not the laws of nature. Let b_z be the posterior after an observation with marginal probability $p(z)$. Then

$$\sum_z p(z) D_{\mathrm{KL}}(b_z \| b) = I(\Theta; Z). \tag{9}$$

The identity follows by substituting $b_z(\theta) = p(\theta, z)/p(z)$ into the KL divergence. It measures expected discrimination among the declared hypotheses. It does not assign Canadian dollars, joules or independent proof to a bit of information.

For a differentiable model p_η, the local expansion $D_{\mathrm{KL}}(p_{\eta+d\eta} \| p_\eta) = \frac{1}{2} d\eta^\top F(\eta) d\eta + o(\|d\eta\|^2)$ defines a Fisher-information metric under standard smoothness conditions. A natural-gradient direction $F^{-1}\nabla V$ is meaningful only for the specified objective, coordinates and positive-definite metric. It is not evidence that a corporate mission obeys physical least action.

The supplied Successor Manifold is a space of complete reachable architectures, not simply neural weights. An architecture can change through a better instrument, rights-cleared evidence, a more faithful Gym, a different conventional solver or a lower-burden review process. Those changes can have discontinuous constraints. A smooth geometry is one local analytical lens, not a universal optimizer or a replacement for discrete gates.[1, 2]

13 Recovery physics: conservation, separation and exergy

The retained process account is phenomenological: its normalized variables and yield polynomial are authored. A real recovery mission would additionally require composition, phases, reaction kinetics, heat transfer, transport, equipment efficiencies and calibrated measurements. A fitted scalar response alone cannot establish these mechanisms.

A continuum description balances species inventories through transport and reaction sources. Individual species need not be conserved when reactions interconvert them. If m is the vector of inventories, N a consistently unit-weighted stoichiometric matrix and r reaction rates, then $\dot{m} = Nr + f_{\mathrm{in}} - f_{\mathrm{out}}$. Any conserved weight vector w with $w^\top N = 0$ obeys

$$\frac{d}{dt}(w^\top m) = w^\top (f_{\mathrm{in}} - f_{\mathrm{out}}). \tag{10}$$

Accounting cannot create accepted mass by renaming a transfer or counting recovered feed twice.

For an ideal mixture of n moles at fixed temperature and pressure, complete reversible separation into pure components requires at least

$$W_{\min} = -nRT \sum_i x_i \ln x_i. \tag{11}$$

This follows from the difference between the mixed and unmixed Gibbs free energies. At 300 K, one mole of an equimolar ideal binary mixture requires approximately 1,729 J in this reversible limit. Real impurities, irreversible transport, finite throughput and nonideal chemistry can require much more. The expression is neither the plant's estimated electricity bill nor a demonstration that the authored yield is achievable.

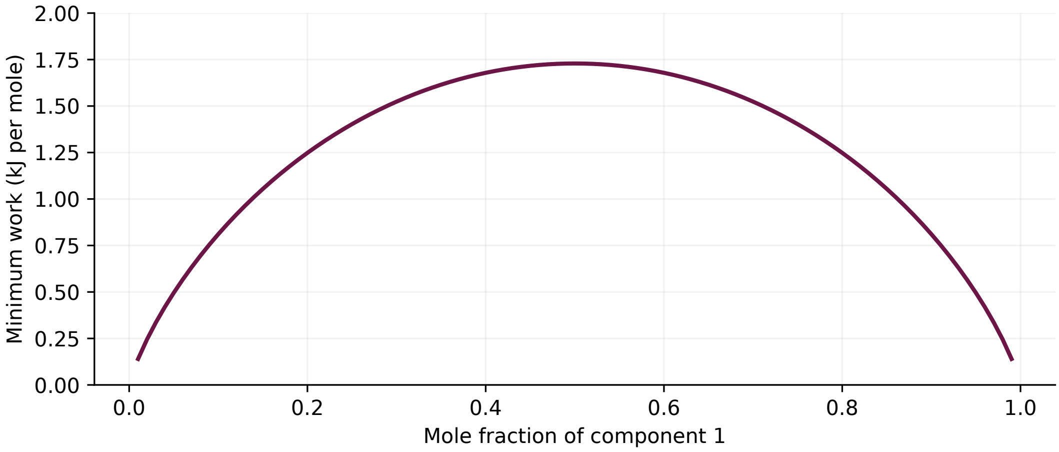

Figure 5 – *Ideal reversible separation work for one mole of a binary mixture at 300 K. The curve is a thermodynamic reference calculation, not a recovered-material dataset. SI constants follow the BIPM reference.[11]*

The owner should ask whether accepted mass, quality, specific work and heat rejection improve together. Increasing throughput by exporting unaccounted entropy, scrap or downstream risk is not a verified capability gain.

14 Physical feedback and institutional free energy

A physical port-Hamiltonian model[12] has $\dot{x} = (J - R)\nabla H + Bu$, with $J = -J^\top$ and $R \succeq 0$. Direct multiplication gives

$$\dot{H} = -\nabla H^\top R \nabla H + y^\top u, \qquad y = B^\top \nabla H. \tag{12}$$

Thus dissipation cannot be negative in that declared model. A proposed controller must still specify the state, storage function, coupling, uncertainty and implementation. No such physical controller is activated by this publication.

At mesoscopic scale, transition currents $J_{ij} = p_i k_{ij} - p_j k_{ji}$ yield a nonnegative entropy-production rate

$$\dot{S}_{\text{tot}} = \frac{k_B}{2} \sum_{i,j} J_{ij} \ln \frac{p_i k_{ij}}{p_j k_{ji}} \geq 0 \tag{13}$$

for a finite Markov process with reciprocal positive rates. Each paired term uses $(x-y)\ln(x/y) \geq 0$. This is *total* entropy production, not environmental entropy alone. With $\dot{p}_i = -\sum_j J_{ij}$,

$$\dot{S}_{\text{sys}} = -k_B \sum_i \dot{p}_i \ln p_i, \tag{14}$$

$$\dot{S}_{\text{med}} = \frac{k_B}{2} \sum_{i,j} J_{ij} \ln \frac{k_{ij}}{k_{ji}}, \qquad \dot{S}_{\text{tot}} = \dot{S}_{\text{sys}} + \dot{S}_{\text{med}}. \tag{15}$$

A symmetric two-state chain with $p = (0.9, 0.1)$ and unit reciprocal rates has $\dot{S}_{\text{med}} = 0$ but $\dot{S}_{\text{tot}} = 0.8 k_B \ln 9 > 0$. The new executable counterexample checks this decomposition. Local detailed balance and an appropriate state/coarse-graining convention are required to interpret the medium term thermodynamically; arbitrary software transitions are not automatically heat flows.[13]

Information can enter a physical feedback bound, but only with a physical protocol. For an initially equilibrated system and a measurement-feedback experiment satisfying the generalized fluctuation relation,

$$\left\langle e^{-\beta(W-\Delta F)-I}\right\rangle = 1 \quad \Longrightarrow \quad \langle W\rangle \geq \langle \Delta F\rangle - k_B T\langle I\rangle. \tag{16}$$

The implication is Jensen's inequality; the equality's assumptions belong to the thermodynamic protocol.[14] There is no conversion factor taking mission NPV into extractable thermodynamic work.

A separate mathematical identity does connect entropy-regularized choice to a Gibbs distribution. For finite feasible actions, potential Φ and $\tau > 0$,

$$-\mathbb{E}_\rho \Phi - \tau H(\rho) = -\tau \log Z + \tau D_{\mathrm{KL}}(\rho \| \pi), \quad \pi(a) = e^{\Phi(a)/\tau}/Z. \tag{17}$$

Here τ has the units of the potential, not necessarily kelvin. Asynchronous logit dynamics have this stationary law in a finite exact-potential game under the appropriate update rule.[15] Neither the retained planner nor the new cooperative game is asserted to be such a strategic dynamical system.

The institution that can say no

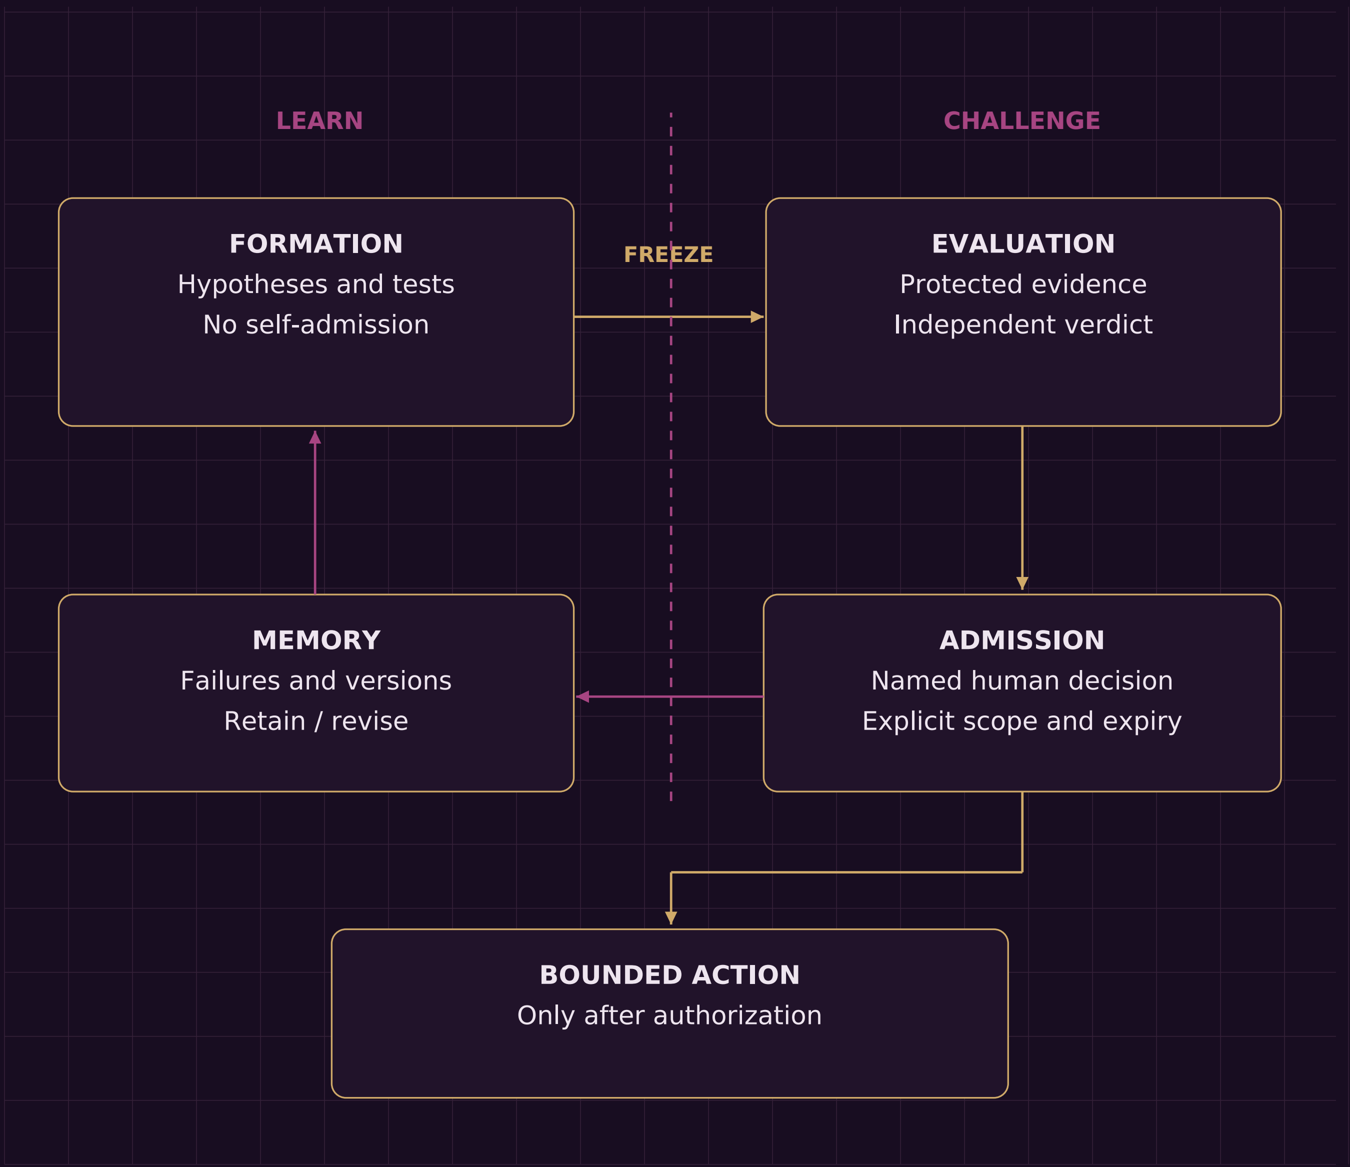

Target architecture — no authority granted here.

A useful result survives a serious alternative, a fair accounting boundary and a review it does not control.

15 The game of access to evidence

The new cooperative calculation contains an owner and five evidence providers. Assay prices are already deducted inside the project value. A coalition without the owner has zero project value: it lacks the capital, asset rights and decision. A coalition with the owner may use the subset of tests supplied by its members, under the unchanged three-test and cash limits.

Let $v(S)$ be the optimal expected NPV with that access. All 32 test menus are solved, producing 64 owner/provider coalitions. The Shapley allocation is

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} [v(S \cup \{i\}) - v(S)]. \tag{18}$$

This is a standard marginal-contribution allocation, not a price mechanism or a guarantee of stable agreement.[16]

Proposition 8 (An owner-essential monotone game has a core)**.** If tests are optional, $v(S)$ is monotone in access. If $v(S) = 0$ without owner 0, $v(N) \geq 0$, and transfers are unrestricted, then allocating all $v(N)$ to the owner is in the core.

Proof. A coalition without the owner gets zero, matching its value. A coalition containing the owner receives $v(N) \geq v(S)$. The allocation is efficient. Optional access preserves existing choices and hence monotonicity. □

That existence result does not imply desirable bargaining power or realistic provider participation at unknown prices. It is conditional on paid assay contracts and the declared transferable surplus.

The computed Shapley allocation assigns CA\$57.729 million to the owner, and approximately 7.295, 6.628, 7.351, 15.493 and 10.168 million to the five evidence services. Twelve proper coalitions can block it, with a maximum deficit of CA\$11.917 million. A core witness exists; one LP solution allocates 103.664 million to the owner and 0.999 million to recovery A. It is not uniquely fair or a recommended commercial contract.

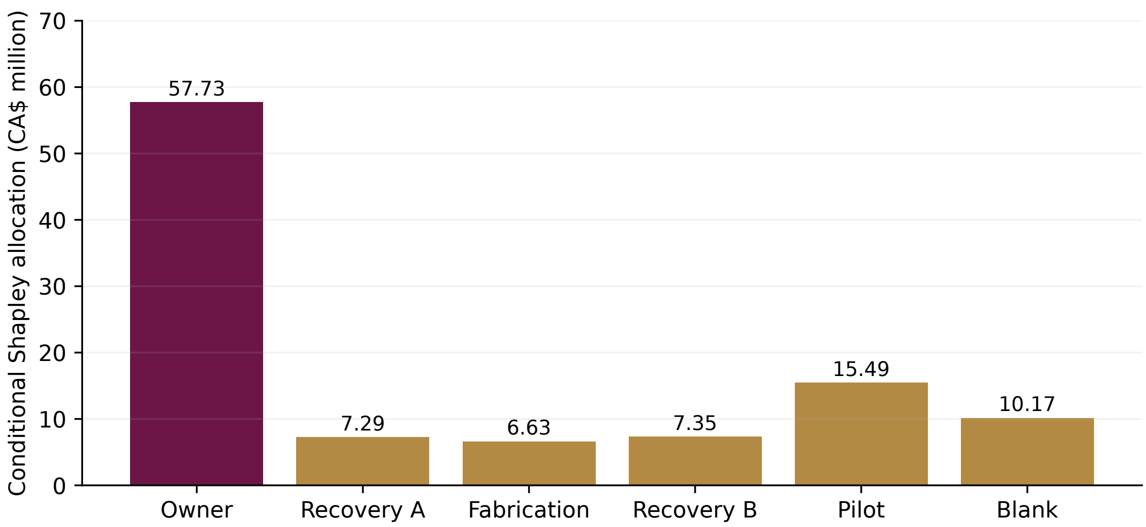

Figure 6 – *New conditional residual-surplus allocation after assay costs. Provider labels indicate access to the respective test, not an actual payment or an independent organization.*

The lesson is institutional: optional information can be complementary, redundant or substitutable. Crediting every provider with the full value of the same project manufactures money.

16 Incentives for useful measurements, not favourable verdicts

A claimant can influence the apparent case by designing a signal, selecting what to disclose or presenting correlated witnesses as independent. Bayesian information design formalizes why even truthful signals can be chosen to influence a receiver.[17] The appropriate response is a constituted observation channel, declared negative outcomes and separated custody—not an assumption that every confident report is false.

Proposition 9 (Truthful reporting under a proper quadratic score)**.** For binary outcome Y with true conditional probability p, expected score $\mathbb{E}[-(r-Y)^2] = -(r-p)^2 - p(1-p)$ is uniquely maximized at report $r = p$.

This encourages honest belief reporting only when the scored outcome is externally grounded and no other payoff reverses the incentive. It does not encourage costly investigation by itself.[18]

Proposition 10 (The price of information acquisition)**.** If the investigator observes signal Z, reports truthfully, receives score multiplier k and incurs cost c, the expected incremental score relative to reporting the prior is $k\,\mathrm{Var}(\mathbb{E}[Y \mid Z])$. Investigation is privately worthwhile precisely when that amount exceeds c, absent other payoffs.

Proof. The truthful expected quadratic score is minus conditional Bernoulli variance. Apply the law of total variance to the difference before and after the signal, then subtract c. □

Compensation contingent only on a candidate passing a test creates a different incentive. A proof office should purchase calibrated, complete measurements rather than a particular promotion outcome.

Proposition 11 (A scoped audit deterrence condition)**.** Suppose false reporting creates gain g if undetected, audits occur with probability a, detection probability is d, an enforceable loss is F, and an honest report faces false-positive probability f. Under the specified payoff model, cheating is deterred when

$$a\,[d(g+F) - fF] \geq g. \tag{19}$$

Proof. Honest incremental utility is $-afF$. False reporting yields $(1-ad)g - adF$. Compare the two. If the denominator is nonpositive or the required a exceeds one, this contract cannot deter the stipulated gain. □

The formula does not estimate enforceability, detection or human preferences. Separation of formation, proof and admission remains necessary even when the algebra appears favourable.

17 Mission Alpha and successive opportunities to be wrong

An authored no-build baseline is not the strongest available alternative. For a frozen mission and common utility basis, a proposed conservative advantage is

$$\alpha^- = \min_{b\in\mathcal{B}} \mathrm{LCB}[U(S) - U(b)] - C - R - H - I - D - \Gamma_{\mathrm{Gym}}, \tag{20}$$

where the final term is a reserve for Gym-to-reality mismatch. Deductions must be incremental, horizon-consistent and counted once. The source papers use related templates, not a universal accounting standard.[1–3]

The central informed OLS comparator receives the same permitted training observations and full contingent planner, and ties the candidate. Subtracting the separate CA\$12 million allowance yields CA\$−12 million relative to that comparator. Against information per dollar it yields CA$−4.087 million; against one-step research, CA$38.791 million. Missing field costs or basis-risk reserves cannot be treated as evidence of zero burden.

Repeated successor searches also require statistical discipline. As an illustrative future protocol, allocate $\delta_{gjt} = 36\delta/(\pi^4 g^2 t^2 K_g)$ to generation g, comparator j among K_g, and inspection time t. Their total is δ. Condition on each generation's formation history: the candidate, finite comparator set K_g, and payoff definition must be fixed before new evaluation draws. For differences in a declared range of width R_0, with fixed conditional mean and independence within that fresh evaluation stream, a one-sided Hoeffding radius is $R_0\sqrt{\log(1/\delta_{gjt})/(2t)}$. A union bound then controls the family under those assumptions. More efficient time-uniform methods exist.[19]

The retained exact synthetic expectations need no Monte Carlo confidence interval. This proposed statistical rule is for a future evidence-generating process. Adaptive reuse of exposed data, clustered failures or a changed comparator can invalidate its assumptions. A field claim must identify that process and its real control arrangements; hashes cannot create independence.

Build the next feasible world

Accepted value opens a new mission. It cannot waive conservation, commissioning, replacement or proof.

18 Capital comes back before surplus compounds

The separate favorable-state continuation is not the average outcome of the decision tree. Recovery at intensity 1.5, with suitable-state contamination 0.10, produces 139,860 tonnes from 180,000 tonnes of feed, leaving 40,140 tonnes of residue. Electricity is 236,250 MWh. External sales are CA$195.804 million and first operating net is CA$107.026 million before tax and financing.

Equipment costs CA$360 million. Charging the full CA$6.5 million assay ceiling and CA$12 million programme allowance creates CA$378.5 million initial funding. Later operating net declines by the assumed 1% annual retention factor. No cash arrives during commissioning. Contributed capital is repaid before allocations to science, distributions or replacement.

Twelve-year favorable-state account	CA$ millions
Initial equipment, assay ceiling and programme funding	378.500
Project NPV after those initial costs	213.945
Contributed capital returned	378.500
Science allocation after capital return	128.972
Surplus distributions after capital return	128.972
Replacement reserve after capital return	386.917

Returned capital is not income. The unsuitable-state route has CA$−152.862 million NPV under the same spending assumptions. Both records remain inspectable; their difference is why investigation matters.

The source value ladder proceeds through identified, modeled, validated, authorized, implemented, realized, collected or conservatively attributable, and reinvestable value.[1] ANABASIS's monetary ledger remains in the modeled state. It does not execute a transfer, approve a tax treatment or assign operational authority.

19 Capacity follows the binding constraint

A separate throughput account supplies 135,000 tonnes of recovered material. Fabrication accepts 110,000 tonnes and yields 101,200 tonnes. Inspection can qualify only 80,000 tonnes. Thus 21,200 tonnes await inspection, while 25,000 tonnes of recovered feed remain unused.

The next investment is not automatically a larger recovery plant. Raising inspection to 101,200 tonnes removes that bottleneck. Raising it further to 140,000 tonnes produces no additional accepted output in the stated network. Prices and engineering costs are intentionally absent from this isolated demonstration: a throughput improvement is not yet an investment return.

For a general material network, let A encode input requirements and x activities. A feasible expansion must satisfy material, energy, inspection and resource inequalities, as well as cash timing. The shadow value of a binding constraint can help screen a next mission, but a linear-programme dual price is conditional on its local technology, prices and feasible region. It is not a permanent market price.

The next-successor rule

Fund a capability that changes the limiting feasible set, not merely a component that looks impressive in isolation. Preserve the current method until the descendant has independently earned the affected claims.

The source definition of recursive institutional improvement requires fresh-work gains in the institution's ability to form, prove or govern a better successor without weaker proof or governance.[1, 2] Reusing a template or a previous equation establishes recurrence. A version number alone establishes neither learning nor a compounding institution.

20 A growing estate must also replace itself

Let $M(t)$ be useful installed stock, $B(t)$ accepted installations, growth $g \geq 0$, and lifetime $L > 0$ deterministic. For a balanced trajectory $B(t) \propto e^{gt}$,

$$M(t) = \int_{t-L}^{t} B(s)ds, \qquad \frac{B(t)}{M(t)} = \frac{g}{1 - e^{-gL}}, \tag{21}$$

with limit $1/L$ as $g \to 0$. Retirement is $R = Be^{-gL}$, and $B - R = gM$.

Proposition 12 (Virgin-input lower bound)**.** If manufacturing yield is $0 < y \leq 1$, recovered retirement fraction is $0 \leq r \leq 1$, and only retirements supply secondary feed (no manufacturing-scrap return or inventory draw), then gross process feed is B/y and required virgin input is

$$V = \frac{B}{y} - rR = gM + \left(\frac{1}{y} - 1\right)B + (1 - r)R \geq gM. \tag{22}$$

Proof. Substitute $B - R = gM$. The remaining terms are nonnegative under the stated bounds. □

For $M = 10^{22}$ kg, $g = 2\%$ per year, $L = 50$ years, $y = 0.9$ and $r = 0.85$, accepted installations are 3.164×10^{20} kg/year, retirements 1.164×10^{20}, gross feed 3.516×10^{20} and virgin input 2.526×10^{20} kg/year. These are ideal steady-flow calculations, not forecasts of manufacturing capability.

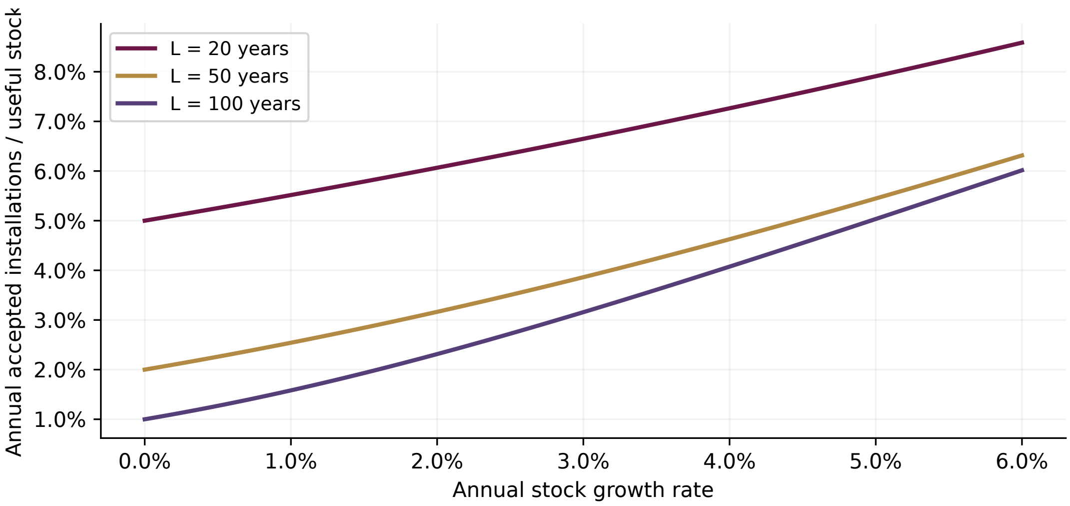

Figure 7 – *Gross installation burden relative to useful stock. More efficient recycling does not remove manufacturing losses or the material required for net growth.*

A finite history has transient commissioning and retirement waves; a distributed lifetime needs an age-structured survival model. Every extension adds evidence obligations, not just another growth parameter.

21 A stellar-scale horizon with a physical denominator

The prospective estate connects science, material recovery, fabrication, qualification, energy capture, computation and repair. It is a network of bounded mission institutions, not a controller whose authority increases merely because its ambitions do.

For an illustrative useful-power target $P_u = 10^{26}$ W, conversion efficiency $\eta = 0.30$ and an availability factor $a = 0.97$, the nameplate capture requirement is

$$P_c = \frac{P_u}{\eta a} = 3.436 \times 10^{26}\ \mathrm{W}. \tag{23}$$

This is 89.77% of the IAU nominal solar luminosity 3.828×10^{26} W. The nominal irradiance at one astronomical unit is 1361 W/m^2; these are exact conversion conventions, not instantaneous measurements of the Sun.[20]

Ideal unshadowed collector area is $P_c/1361$. If all nameplate-captured power must ultimately be rejected inside the estate, a conservative effective emitting-area account is

$$A_r = \frac{P_c}{\epsilon\sigma(T_r^4 - T_s^4)}. \tag{24}$$

At $T_r = 600$ K, $T_s = 3$ K and $\epsilon = 0.9$, it requires approximately 5.196×10^{22} m^2. Emitting area is not planform area; geometry, view factors, solar absorption, mutual heating and degraded surfaces can change the result. NASA's thermal-control reference treats such system-level energy balances; it does not validate this hypothetical estate.[21]

At 20% efficiency the same target and availability exceed one nominal star's budget. No economic curve can repair that arithmetic. Structures, transport, repair facilities, fabrication energy, rare materials, demand and governance remain unsolved mission requirements in this scenario.

Finite propagation also constrains institutions. For distance d, a round-trip observation-and-response loop needs at least $2d/c$ before processing and actuation. At one astronomical unit this is approximately 998 seconds. Fast local protective functions therefore need qualified local control; remote strategic supervision cannot replace them. No stellar arrival date or valuation is extrapolated from the twelve-year recovery project.

From a first mission to a successor

One decision. One owner. A fair comparison. Explicit evidence gaps. A bounded next experiment.

CONCEPT ART OR SCHEMATIC · NOT AN INSTALLATION OR EVIDENCE

22 Using SUCCESSOR Ω generally

The most useful first mission is repeated, consequential, observable and bounded. It has an accountable principal, an identifiable incumbent and a credible alternative. A free-text wish to transform a business is not yet a numerical contract.

Start with the decision and its outcome. Record the available information at decision time, action limits, unacceptable failures, time horizon, costs and rights. Separate supplied constraints from mechanisms to be learned. Name what the chosen adapter cannot represent. Before fitting, define development and challenge evidence and the comparison that will decide whether the result is useful.

In the included study, `FIRST_MISSION.html` prepares an exportable draft. It does not train a model, authenticate evidence or issue permission. `MISSION_COMMAND.html` evaluates the finite ANABASIS example. Its arbitrary capital numbers must never be imported into a real mission as facts.

The practical sequence is constitute, inspect evidence, form, freeze, challenge, compare, decide, export and rehearse restoration. A completed local run may correctly fail or remain inconclusive. Keep the exact candidate and negative result instead of changing the test until it looks favourable.

The separate guide in this publication includes an offline reading route, a runnable native v10.2 example, a v11 reference route, an owner review template and a handoff checklist. The retained general Field Guide provides additional examples and their original limitations.

23 Version-specific use and truthful portability

The scientific study ran on v10.2. The new v11.0.0 product adds a persistent local evidence-to-successor workflow, explicit comparators and exposure records. It is a separate release, not a retrospective upgrade to the study's evidence.[22]

For this publication's bundled research route, run the commands from the study root using the prepared dependencies. A fresh output folder avoids overwriting historical evidence. The native numerical tutorial should produce `CANDIDATE`, then `LOCAL_PASS`, while authority remains false.

For a separately installed v11 release, its guide uses `SUCCESSOR.py doctor`, `init --workspace ...`, and `serve --workspace ...`. The loopback session link is private; the local service is not designed to be exposed through a public tunnel. The v11 executable is not bundled in this publication. Its actual guide is included as a reference, so incompatible command sets are not mixed.

Portability is demonstrated only to the extent tested. A ZIP and hash list provide inspectable bytes and local integrity. They do not establish legal ownership, an authenticated publisher, independent replication or customer-controlled production. Rights-aware export must preserve evidence restrictions and actual exposure history. A restored artifact does not restore expired proof or operating permission.[3]

24 The complete handoff and admission boundary

The 21-job map identifies what formation, benchmarking, rights, security, economics, resilience and independent acceptance require. In this case, source-linked records exist for local mission formation, learned models, finite decisions, failure taxonomy and replay. External apparatus, protected institutional custody, actual commercial costs and accountable mission admission remain requirements.

Review question	What would satisfy it
Is the mission observable?	Qualified measurements, valid event identities, uncertainty and rights.
Is the comparator credible?	A current procurable alternative under the same inputs and cost basis.
Is the evidence fresh?	Protected cases not used for fitting, selection or repeated post-hoc adaptation.
Does the full system improve?	Comparative decision outcomes, critical failures, burden and all hard gates.
Can it be operated and replaced?	Named authority, scoped permissions, expiry, monitoring, rollback and recovery drills.

The admitted state called Specialist ASI in the source is complete-system superiority within a constituted mission. SUCCESSOR Ω admission additionally involves continuity, rights, identity and separately granted scope.[1, 2] Neither designation follows from a generated title, a signed local file or the prominence of the cover art.

The useful near-term outcome is an institution able to submit a credible review request: exact frozen composition, complete denominator, declared limitations, competitor evidence and an operationally meaningful stop condition. The owner must retain the ability to reject a favourite candidate and continue with a better-supported predecessor or supplier.

25 Limitations and the expansion thesis

Industrial parameters, observations, prices and hidden states are authored. The central scenario was chosen to expose option value and shared-channel risk, not to represent a population of industrial projects. Pilot laws and topology are supplied; evidence remains under common author custody. Bounded numerical search is not unrestricted scientific discovery.

The loss gate concerns one pre-assay NPV threshold, not complete risk preference or engineering safety. Historical pruning keeps a potentially wasteful test sequence. Appendix I separately reoptimizes nominal value under a specified persistent-channel family; it establishes neither worst-case expected-value optimality nor field validity. The game assumes an essential owner and transferable surplus after assay costs. Thermodynamic identities apply to specified physical systems, not to mission narratives.

The informed conventional comparator ties the original decision. Realized customer Alpha remains unmeasured and credited zero. Neither locally reproduced calculations nor a new version establishes independent superiority, operating authority or institutional compounding.

The constructive asset is nevertheless concrete: a reusable way to constitute a mission, purchase decision-changing evidence, learn executable hypotheses, account for consequences, challenge alternatives and preserve failures. The owner retains that method when a particular model is replaced.

The expansion institution

A greater productive horizon begins with a decision that survives its evidence. Preserve the institution that can keep making that distinction when the world, its models and its suppliers change.

A Notation and measurement boundaries

Symbol	Meaning and unit
u, z	Normalized intensity and contamination; dimensionless.
$y(u, z)$	Recovered mass fraction; not a remaining-life or causal model.
$\theta = (r, f, b)$	Eight-state suitability and common-bias configuration.
b_h, L_a	Posterior distribution and observation channel; probabilities.
U_d, c_a, K_d	Project NPV, assay cost and construction capex; tables use CA$ millions unless stated.
$V(h)$	Optimal remaining conditional value; past assay costs are sunk.
$\mathcal{D}(h)$	Actions passing the frozen capex and posterior loss conditions.
δ	Statistical error allocation in a proposed future protocol, not the synthetic world's failure probability.
F, T, I, G, S	Feed, internal transfer, imported feed, finished mass and fabrication scrap; tonnes per year.
$M, B, R, V_{\text{virgin}}$	Installed stock and accepted, retired and virgin mass flows; kg and kg/year.
g, L, y, r	Growth per year, lifetime in years, manufacturing yield and recovery fraction.
P_u, P_c, η, a	Useful and captured power in watts, efficiency and availability.
H, Φ, τ	Physical stored energy, game potential and regularization scale. Their units are not interchangeable.

Assay predictions, finite expected returns and physical measurements are different objects. Inputs described as normalized must be mapped to calibrated physical units before real engineering use. The BIPM reference controls SI conventions; commercial cash units remain separate.[11]

B Exact terminal account and loss definition

For terminal design d, operating cash before tax and financing is external recovered and finished sales less purchased feed, electricity, fixed operating cost and disposal. Capex is paid at time zero. For $H = 12$, commissioning $\ell = 2$, retention $\gamma = 0.99$ and discount $r = 0.08$,

$$U_d = -K_d + \sum_{t=\ell+1}^{H} \frac{O_d \gamma^{t-\ell-1}}{(1+r)^t}. \tag{25}$$

The risk event uses this U_d, not $U_d - \sum c_a - 12$. The optimization charges assay costs once through the Bellman recursion. The paper's complete-burden diagnostics subtract the separate programme allowance afterward, without pretending the gate evaluated that different loss definition.

No-build is not cost-free after research: its project value is zero, but assay fees already purchased remain lost. A failed project can also require further cash; the favorable capital example cannot justify allocating hypothetical future receipts.

C Posterior sufficiency, ordering and scope

With static hidden state and conditionally independent channels, the likelihood product in (2) is invariant to observation order. Therefore a canonical sorted history is sufficient for the retained recursion, provided spent cost and remaining test availability are recovered from that history. This does not hold automatically for evolving calibration, instrument wear, learning-by-testing, or observations whose cost and law depend on order.

The policy chooses tests adaptively; it does not condition on the hidden state. The state is sampled once in each local Gym episode and exposed only through permitted observations and the terminal account. Such API separation is not OS-level isolation. An independent evaluator must control the protected cases, scorer and relevant observation process outside claimant write access.

The likelihood calibration counterexample uses two tests with the same parameters but distinct purchased identities. They are conditionally independent given the bias. Replacing their joint law with the product of their marginals destroys exactly the dependence that repeated positives should reveal.

D Finite robust guard: complete interpretation

The new guard tests four alternative biased-pilot positive probabilities while retaining the original priors, physical responses, budgets and test sequence. At every terminal branch it evaluates the same severe-loss indicator against each resulting posterior. A violating investment is replaced with no-build. Its formal guarantee is restricted to those models.

The published 38-point diagnostic is a sensitivity study across probabilities 0.025 through 0.95. It does not establish a confidence interval, a calibrated distribution over model error or a guarantee under all future common-mode failures. The conservative value decrease is a price of this particular guard, not the optimal cost of robustness. Edition 1.3 adds that separate nominal-objective reoptimization in Appendix I, with a declared interval family and no reclassification of these historical results.

E Coalition certificate and its proper use

The new information game has $2^6 = 64$ coalitions, 32 of which contain the essential owner. Test fees are part of every menu's already computed expected value. Marginal contributions are verified by all $6! = 720$ player orderings, not merely by checking that allocations sum to the total.

A core certificate requires $\sum_i x_i = v(N)$ and $\sum_{i \in S} x_i \geq v(S)$ for every proper coalition. The least-core programme minimizes nonnegative ϵ subject to the relaxed inequalities. Its value is zero in this owner-essential monotone game. The Shapley vector is efficient but not in the core; efficiency alone is not coalition stability.

The reported allocation is one feasible witness, not an optimal social contract or a legal allocation of rights. Changing provider outside options, assay prices, ownership, exclusivity or nontransferable preferences changes the game. The owner-essential assumption is explicit precisely because omitting capital ownership could falsely let several coalitions each claim the same project's entire value.

F How the source documents are used

Source location, PDF pages	Application and boundary
Navigator v5.0.0, 31–34	SEIZE and distinct decision functions; Elevate remains a review request.
Navigator, 45–46	Value-state ladder and fresh-work recursive improvement; no projected cash authorization.
Gym v2.1.0, 16–21	Controlled mission world, role observations and five Gym families.
Gym, 46–48	Environment revision, negative capability graph and non-inherited authority.
Gym, 72–75	Canonical state machine, 21 bounded jobs and reference records.
Corporate white paper, 21	Frozen serving branch versus separate learning successor; no unreviewed cutover.
Corporate white paper, 50–54	Maturity and claim-specific evidence; architecture is not an end-to-end deployment claim.

Source terminology is preserved without turning physical metaphors into laws of institutional motion. The Gym paper's Alpha template explicitly includes basis risk. Decision functions, D0–D4 phase gates and A0–A4 authority are separate constructs. No token ownership check, on-chain job or wallet approval is required for this research archive.

G Reproduction and provenance

The retained study lives under `study/ANABASIS_OMEGA_1_0/`. Its original manifest and evidence are preserved. Run its verifier and tests from that directory. Its replay command creates a new workspace; a new folder must be used for each run.

```
python3 VERIFY_PACKAGE.py
python3 REPRODUCE.py --out /path/to/new/reproduction
python3 -m pytest tests -q
```

From the publication root, the new study uses the retained library through a relative path:

```
python3 research/frontier.py --out var/new_publication_run
python3 -m pytest research/test_frontier.py -q
python3 VERIFY_PUBLICATION.py
```

The documented authoring environment is recorded in `assurance/VALIDATION.json`. NumPy is required for the retained numerical engine; SciPy supplies the new linear-programme witness; pytest runs the tests. Dependencies are not bundled. The source does not request external services, operate equipment or grant authority.

A new replay reproduced 69 retained result files byte-for-byte. The retained suite passed 68 tests and the publication-side suite passed 41. New outputs were separately regenerated for comparison. PDF rendering, links, layout and archive integrity have their own checks. These checks establish local repeatability and package consistency, not independent scientific confirmation.

H What to do next

The next engagement is one bounded customer mission with qualified measurements and a credible procurable alternative. The owner should first determine whether the pilot fails under the same conditions as the coupons. A repeated observation is not an independent witness merely because it has a different filename.

Use `templates/MISSION_REVIEW.md` to prepare the handoff below. The template records a request and its gaps; completing it does not authorize expenditure, certify evidence or admit a successor.

Review object	Required next evidence
Decision and owner	One consequential choice, named principal, horizon, alternatives and explicit exclusions.
Measurement world	Qualified instrument, event identity, shared failure causes, calibration history and a falsifier for the proposed channel.
Frozen comparison	Exact candidate and credible alternative under the same information, costs and decision constraints.
Evidence custody	Development separated from protected challenge; prior exposure, rights, retention and permitted exports recorded.
Economics	Equipment, assays, human review, integration, downtime, unwind and reserves on one declared time basis.
Decision receipt	Actual inputs, model and Gym identity, full outcome denominator, adverse cases and uncertainty.
Operating readiness	Scope, expiry, stop conditions, recovery target and qualified controls before external effects.
Succession request	What changed, what is reused, what failed, and which affected claims require new proof.

Execute formation and independent review separately. Retain inferior and inconclusive descendants. Establish accepted or conservatively attributable cash before allocating it to science or productive assets; an expected NPV is not a bank balance.

The first useful outcome

One reviewable decision, an honest comparison and a priced next observation are more valuable than an unbounded automation mandate. New physical capability may expand what the estate can attempt; it does not expand a model's authority by implication.

I From pruning decisions to redesigning research

EDITION 1.3 / NEW EXECUTED EXTENSION

The earlier safeguard leaves the original investigation in place and cancels investment at unsafe leaves. That preserves a scoped risk condition, but may still buy observations whose former investment outcome is no longer available. The owner should not assume that a protected decision has also become an efficient research plan.

The new calculation keeps the original physical reference, priors, payoff table, two-year commissioning, twelve-year horizon, budget, catalogue and pre-assay severe-loss event. It changes only the terminal admissibility rule and then recomputes the whole test policy. The nominal observation model remains the expected-value objective. The risk model set contains that nominal law *plus* the persistent family in which a biased instrument gives a positive pilot with probability $p \in [0.12, 0.84]$. Every other channel is unchanged.

A different question, with a narrower answer

Maximize nominal project NPV after assay costs, subject to the posterior loss gate under every declared channel. This is not maximization of worst-case expected NPV; nature does not choose a new adverse model at each history. No probability distribution over the model interval is invented.

	Original	Pruned	Replanned
Nominal NPV after assays, CA$ M	104.664	78.345	86.901
Expected assay cost, CA$ M	2.859	2.859	1.054
Probability of no-build	47.86%	60.22%	55.53%
Worst nominal leaf loss risk	7.72%	7.72%	9.68%

The replanned policy recovers CA$8.556 million relative to pruning, including CA$1.806 million less expected assay expenditure. Those amounts are not additive: the fee saving is already reflected in the NPV difference. Replanning uses more of the frozen 10% risk budget; it is not uniformly less risky than pruning. All amounts are conditional research calculations, not receipts or product Alpha.

What the owner actually does differently

Start with the low-cost recovery coupon B. A negative result stops immediately. A positive result leads to the fabrication coupon. A negative fabrication result supports recovery at intensity 1.5; a positive result leads to the calibration blank. A negative blank supports recovery; a positive blank stops. The new plan never purchases the joint pilot.

The fabrication observation is useful even though the chosen asset is recovery-only: through the shared bias variable, it changes the interpretation of the earlier recovery observation. This value depends on the declared prior and channel relationships, not on a universal rule that one should test an unrelated process. The four reached leaves, their probabilities, costs and decision identities are exported by `research/replan.py`.

Evidence branch	Probability	Total assay fees, CA$ M
B negative: no-build	29.160%	0.450
B positive; fabrication negative: recovery	29.962%	1.100
B and fabrication positive; blank negative: recovery	14.509%	1.450
B, fabrication and blank positive: no-build	26.369%	1.450

An interval result rather than a grid assertion

Let L_p equal the nominal likelihood matrix except that the biased-state pilot column is p. A test can be purchased at most once. For any history h and terminal decision d, define

$$q_{h,d}(p) = \Pr_{L_p}\{U_d < -100 \mid h\}. \tag{26}$$

Proposition 13 (Endpoint reduction for a once-used affine channel)**.** For a fixed prior, fixed terminal loss event and fixed other channels, assume all considered history probabilities are positive on $[p_-, p_+]$. If the uncertain channel is affine in p and appears at most once in h, then

$$\sup_{p\in[p_-,p_+]} q_{h,d}(p) = \max\{q_{h,d}(p_-), q_{h,d}(p_+)\}. \tag{27}$$

Proof. Each state's unnormalized history weight is affine in p: it contains at most one factor p or $1-p$, or no uncertain factor. Both the loss-weighted numerator and total denominator are therefore affine. Write their ratio $(a+bp)/(c+dp)$. Its derivative is $(bc-ad)/(c+dp)^2$, of constant sign because the denominator is positive. Hence the maximum is attained at an endpoint, including the constant case. □

The nominal model is separately included because its biased-state pilot probabilities depend on suitability and need not correspond to one constant p. The interval proof fails without further work for repeated uncertain assays, changing priors, nonlinear channel families, hidden state evolution, a changed loss definition or new candidate-dependent payoffs. The 145-point numerical sweep is an additional diagnostic, not the proof.

Let $\mathcal{D}_*(h)$ contain only terminal choices satisfying capex and the risk gate for the nominal law and both endpoint laws. Define

$$V_*(h) = \max\left\{\max_{d\in\mathcal{D}_*(h)} b_h^\top U_d,\ \max_{a\in\mathcal{T}(h)}\left[-c_a + \sum_o \Pr_0(o \mid h,a)V_*(h,a,o)\right]\right\}. \tag{28}$$

Here $\Pr_0$ and b_h remain nominal. No-build remains available, and previous assay costs are sunk, not subtracted twice.

Proposition 14 (Optimal nominal value under persistent-channel leaf gates)**.** Under the fixed finite problem above, backward induction computes an optimal deterministic contingent policy among those whose reached terminal decisions satisfy the declared channel-family gate. Its nominal value is at least that of any feasible pruning of the original policy in the same budget and action catalogue.

Proof. At the maximum depth, the best admissible terminal choice is optimal. At every earlier history, each permitted test has a fixed nominal observation distribution and leads to optimal admissible continuations by induction. Taking the maximum of testing and stopping is optimal. The pruned policy belongs to that feasible deterministic class, so its value cannot exceed the new maximum. Endpoint reduction establishes membership for the interval family; it does not change the expectation law. □

This is distinct from a minimax Bellman recursion, whose dynamic consistency can depend on rectangularity of the ambiguity set.[9] Randomized terminal action mixtures, generalized test designs and real-world observation uncertainty are outside the evaluated policy class.

Numerical checks, comparative claims and next work

A separately written reference solver reconstructs cash values and uses exact rational arithmetic for the stated channel probabilities and posterior gates. It checks all 131 reachable history states, not only the chosen trajectory. Every action matches; the largest value residual is less than 6×10^{-14} CA$ million and the largest payoff-table residual is less than 1.2×10^{-13} CA$ million. Cash and optimization values still use floating-point arithmetic: this is not a machine-checked proof assistant certificate or independent custody.

An informed conventional predictor can also use this new planner. No exclusive-model advantage is created by supplying only the candidate with better planning logic. Formation, integration, protection of evidence, actual commercial costs and the Gym-to-reality reserve still require measurement. The new CA$8.556 million comparison is *replanning versus pruning*, not a gain above the strongest procurable complete alternative.

The next useful mission

Qualify the observation channel before funding a plant. Record whether the blank, coupons and pilot share failure modes; preserve the unchanged predecessor; then compare complete methods under the same intervention and evidence budget. The current simulation helps define that question. It does not answer it for a real facility.

Corrections made by the critical review

The physical text now separates total, system and medium entropy production and balances species with reaction sources rather than asserting that each species is conserved. The virgin-input equation explicitly excludes manufacturing-scrap return and inventory draw. The future statistical protocol states that candidates and finite comparator sets must be frozen before the corresponding fresh evaluation stream. These are documented clarifications and a corrected thermodynamic interpretation, not new empirical findings.

The package verifier now rejects empty or malformed manifests, duplicate JSON keys, invalid digests, unsafe paths and symbolic links. Its optional strict mode also rejects untracked files in an untouched extraction. The renewal helper rejects non-finite inputs and non-finite results. Integrity remains unauthenticated: a malicious party replacing both payload and unsigned manifest is outside the guarantee.

```
# From the complete publication root, after dependencies
python research/replan.py --out var/replan_001
python -m pytest research/test_referee.py -q
# Before running tools in a fresh extraction:
python VERIFY_PUBLICATION.py --strict
```

The standalone `DECISION_STUDIO.html` recomputes the finite plan locally. Its scenario export carries the actual parameters and results and always leaves production authority false. A changed input is a new scenario, not a new empirical validation. Source references, historical results and licences remain separate from this publication-side extension.

References

[1] Vincent Boucher. Mission-sovereign successors: Goalos singularity navigator Ω + seize. Technical report, MONTREAL.AI and QUEBEC.AI, 2026. Version 5.0.0, 9 August. Exact supplied PDF retained in the accompanying study.

[2] Vincent Boucher. Gym, specialist asi and successor Ω. Technical report, MONTREAL.AI and QUEBEC.AI, 2026. Version 2.1.0, 10 August. Exact supplied 83-page PDF retained in the accompanying study.

[3] Vincent Boucher. Neural-symbolic successor Ω. Technical report, MONTREAL.AI, 2026. Definitive corporate white paper, 25 August; architectural specification, 59 pages.

[4] MONTREAL.AI. Anabasis Ω: The expansion institution, 2026. Research simulation 1.0. Author-controlled synthetic evidence; unchanged SUCCESSOR Ω v10.2 engine. Source, protocols and records retained in study/ANABASIS_OMEGA_1_0.

[5] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.

[6] Silviu-Marian Udrescu and Max Tegmark. Ai feynman: A physics-inspired method for symbolic regression. *Science Advances*, 6:eaay2631, 2020. doi: 10.1126/sciadv.aay2631.

[7] Hao Tang, Darren Key, and Kevin Ellis. Worldcoder, a model-based llm agent: Building world models by writing code and interacting with the environment, 2024. Version 3; research on executable environment models, not validation of SUCCESSOR Ω.

[8] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.

[9] Garud N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005. doi: 10.1287/moor.1040.0129.

[10] David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272, 1953. doi: 10.1214/aoms/1177729032.

[11] Bureau International des Poids et Mesures. The international system of units, 2026. Ninth edition; official 2026 update.

[12] Arjan van der Schaft. Port-hamiltonian nonlinear systems, 2024. URL https://arxiv.org/abs/2412.19673. Physical modelling and control framework; not an implemented ANABASIS controller.

[13] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75:126001, 2012. doi: 10.1088/0034-4885/75/12/126001.

[14] Takahiro Sagawa and Masahito Ueda. Generalized jarzynski equality under nonequilibrium feedback control. *Physical Review Letters*, 104:090602, 2010. doi: 10.1103/PhysRevLett.104.090602.

[15] Pieter Kleer. Sampling from the gibbs distribution in congestion games. *Mathematics of Operations Research*, 48(4):1846–1870, 2023. doi: 10.1287/moor.2022.1322.

[16] Lloyd S. Shapley. A value for n-person games. Technical Report P-295, RAND Corporation, 1952. URL https://www.rand.org/pubs/papers/P295.html. Published in Contributions to the Theory of Games II, 1953, pp. 307–317.

[17] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615, 2011. doi: 10.1257/aer.101.6.2590.

[18] Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437.

[19] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *Annals of Statistics*, 49(2):1055–1080, 2021. doi: 10.1214/20-AOS1991.

[20] Andrej Prša, Petr Harmanec, Guillermo Torres, et al. Nominal values for selected solar and planetary quantities: Iau 2015 resolution b3. *Astronomical Journal*, 152(2):41, 2016. doi: 10.3847/0004-6256/152/2/41.

[21] NASA. State-of-the-art small spacecraft technology: Thermal control, 2026. URL https://www.nasa.gov/smallsat-institute/sst-soa/thermal-control/. Official engineering reference, accessed 6 September 2026. Not validation of the hypothetical infrastructure in this paper.

[22] MONTREAL.AI. Successor Ω v11.0.0: The evidence-to-successor workflow, 2026. Separate product release. User-guide reference retained; not the implementation on which the ANABASIS trial was run.