TRANSFORMATIONAL INTELLIGENCE ∞

SUCCESSOR Ω

The Institution That Compounds

Manufacturing mission advantage, preserving inventive capability, and financing the physical frontier



The enduring advantage is not a model.
It is the ability to make the next model useful.

Vincent Boucher

MONTREAL.AI • QUEBEC.AI

3 November 2026

A constructive research framework, an exact finite renewal study, and a local authorization-state witness. Not a claim of achieved superintelligence, measured customer profit, or installed physical capacity.

OWN THE INSTITUTION. REPLACE THE MODELS. PRESERVE THE MISSION.

Abstract

A breakthrough can create a result while leaving its owner unable to reproduce, renew, or extend the capability that produced it. We study a different unit of value: a customer-owned mission institution whose methods, evidence, skills, and operating memory can survive changes in models and suppliers without inheriting stale proof or unauthorized powers. We formalize *renewal-adjusted Alpha*: comparative mission advantage over a complete horizon, including construction, unsuccessful search, operation, maintenance, migration, requalification, and residual obligations. The Alpha–Gym–Specialist ASI–SUCCESSOR Ω progression becomes a set of distinct discovery, development, evidence, and authority transitions, not an automatic promotion ladder. An exact authored cash-constrained model exhibits a reversal: a portable capability with lower local operating contribution produces 230 units of horizon cash surplus versus 200 for repeated reconstruction. An independently implemented reachability oracle agrees on 324 optimization comparisons across 162 supplied environments. Only 19 environments favor the additional portability option; 143 tie, and forced retention is often infeasible or inferior. A conventional planner with the same choices ties the result. We derive the crossover condition, specify context-bound renewal and export invariants, and connect sustainable surplus to physical expansion through separately auditable investment and qualification contracts. The contribution is a testable theory of institutional continuity and a reproducible witness of when it matters, not evidence of autonomous discovery, real demand, general ASI, or a forecast of stellar-scale industrialization.

The paper in one sentence

Build advantages worth proving; preserve the machinery that makes them renewable; spend only resources actually available; and make every transition into greater productive capacity earn its own evidence.

Reader's map

Sections 1–3 explain the institution and its four public stages. Sections 4–6 define the evidence boundary, economic objective, and exact results. Sections 7–9 develop renewal, inventive succession, and physical expansion. Sections 10–12 give an actionable industrial research programme, limitations, and the conclusion. The appendices provide proofs, a complete numerical specification, an evidence ledger, and reproduction instructions.

Contents

1 From a successful answer to an enduring institution 3

2 The architecture of ownership 4

3 Alpha, Gym, Specialist ASI, SUCCESSOR Ω 5

4 Preserve the knowledge; renew the warrant 7

5 Renewal-adjusted Alpha 8

6 An exact study of the institution that survives replacement 9

7 What must remain renewable 13

8 The institution must improve more than its inventory 14

9 From productive surplus to a physical frontier 15

10 The next consequential programme: an Industrial Requalification Foundry 17

11 Scope, failure modes, and what would falsify the thesis 19

12 Conclusion: build what can build beyond itself 19

A Formal details and boundary conditions 21

B The numerical record 23

C A claim ledger for the complete programme 24

D Reproduce, inspect, extend 24

1 From a successful answer to an enduring institution

The economically important question is not only whether an intelligence can produce an excellent result. It is whether an owner can reliably obtain further valuable results after the underlying model, operating environment, commercial relationship, or mission requirement changes. A system that wins once but cannot be renewed may be less useful than a slightly weaker system that can preserve, adapt, and re-establish what made its work valuable.

This paper develops *SUCCESSOR Ω* as a proposed institution of mission intelligence. The institution is not a particular model. It is the maintained arrangement through which legitimate knowledge, executable methods, evaluation, and bounded authority serve an enduring purpose. Models are replaceable contributors. They are not the institution's identity, its accumulated evidence, or its right to act.

The founding distinction comes from Transformational Intelligence: the authorized purpose, the current formulation of a problem, and the instrument used to evaluate a result are different objects [4]. A method can change without changing the purpose. A new evaluator requires separate justification. A better-looking score cannot authorize the builder to redefine success. Here we apply that distinction across time: the institution should preserve what remains useful without preserving every assumption, dependency, or permission of the system that originally produced it.

The continuity principle

Continuity is not the refusal to change. It is the ability to change components without losing the mission, legitimate knowledge, or accountability for what has been done.

1.1 What this paper adds

The earlier FLYWHEEL Ω studies how production and verification change feasible economic interactions, including the cash timing that finances subsequent work [1]. SUPER ALPHA connects viable productive arrangements to value capture and successive capability construction [3]. The present paper addresses a different bottleneck: *what survives replacement, and when paying to preserve it is worthwhile*.

We make four concrete additions. First, renewal-adjusted Alpha evaluates an advantage over the lifecycle in which it must remain usable, rather than at one model release. Second, a context-bound succession contract separates portable knowledge from current evidence and authority. Third, an exact finite study identifies when a portability investment beats repeated rebuilding, when it does not, and why a competent ordinary planner can reproduce the result. Fourth, a physical-expansion contract makes the conversion of cash and knowledge into commissioned capacity inspectable rather than automatic.

These are a framework, conditional mathematical results, and newly executed software evidence. The full institutional architecture and industrial programme remain proposed. The public conceptual vocabulary has been synthesized from the author's programme materials; the private working conversations and supplied visual source files are not redistributed. Earlier public papers provide the stable citable foundation. This is a research edition, not a claim of independent peer review.

1.2 The opportunity is more than continuity software

The consequential target is the construction of productive capabilities that organizations cannot currently sustain economically. A laboratory might discover a useful process but repeatedly lose the ability to qualify it after instruments change. An industrial operator might possess excellent models but no portable, testable representation of operating conditions. A group of specialists might solve individual tasks yet fail at the interfaces required for a complete customer outcome.

The institution should therefore search over representations, experimental instruments, interfaces, production sequences, commercial commitments, and coordination arrangements. It may discover that the best move is a conventional method, a smaller system, or no investment. An institution exists to make valuable work possible, not to make itself indispensable.

The long-horizon ambition is correspondingly larger than retaining a chat history. It is to turn successful intelligence into accumulated scientific and productive means, and to use genuinely earned resources to build what the next consequential objective requires. Every step must preserve the difference between a plausible opportunity and a commissioned capability.

2 The architecture of ownership

Represent a mission institution at time t by

$$\mathcal{I}_t = (J, \mathcal{M}, K_t, X_t, H_t, E_t, A_t, L_t), \tag{1}$$

where J is its persistent identity, $\mathcal{M}$ the owner-authorized mission charter, K_t retained knowledge, X_t executable methods and instruments, H_t operating history, E_t current scoped evidence, A_t active authority, and L_t resources and obligations. A model-provider configuration is a replaceable dependency of X_t, not the definition of J.

The distinction between K_t and E_t is essential. The knowledge that a method once passed an evaluation may remain true. The claim that it is qualified *now*, for another environment and mission, may be false. Likewise, knowing how to perform an action does not grant permission to perform it.

2.1 Ownership must have an operational test

Customer ownership is a proposed product requirement, not a legal conclusion established by this manuscript. Its technical acceptance test has four parts: the customer can export the authorized institutional assets; restore them in a separately controlled environment; replace a model or supplier; and recover the relevant bounded behavior under fresh evaluation. Rights to data, software, model weights, and external services must be established separately. Calling a bundle portable does not make a proprietary dependency portable.

The exported assets should include versioned methods, supported input and output contracts, scoped evidence history, provenance, failure conditions, routing procedures, and the real cost of maintaining the bundle. Credentials and active grants are not portable knowledge. A receiving installation begins with no inherited operating authority.

An owner may decide not to preserve an asset whose maintenance exceeds its future value. The institution must support retirement as well as accumulation. Compact, decision-relevant knowledge can be preferable to a larger archive that is expensive, stale, or impossible to retrieve reliably. Historical losses and obligations cannot disappear when operational material is compacted.

2.2 A capability capsule is a hypothesis about reuse

A useful capsule contains an executable artifact, the semantics of its inputs and outputs, prerequisites, dependency identities, prior test evidence, known counterexamples, and a proposed successor use. It does not contain a universal permission to reuse the earlier conclusion. A capsule may be technically intact while inappropriate for its next task.

We distinguish four relationships between campaigns. A predecessor may *finance* a successor, *equip* it with a useful tool, *enable* a construction that would otherwise be inaccessible within matched resources, or *improve discovery* itself. These relationships are not interchangeable. A later result is not evidence of any of them merely because it occurs later.

The strongest institutional asset may be a constructor: a tested method for producing new instruments under changed conditions. But a narrative explaining a successful instrument is not automatically such a constructor. It must generate executable candidates and survive fresh comparisons, including ordinary reuse of the same substantive information.

3 Alpha, Gym, Specialist ASI, SUCCESSOR Ω

The four names describe distinct functions. They need not impose a rigid four-agent implementation, nor require every mission to traverse every stage again. A qualified conventional tool may already be the best route. What cannot be skipped is the evidence or authority required for the actual action.

3.1 Alpha: find the advantage worth building

Alpha begins with a beneficiary and a consequential outcome, not a favorite architecture. Identify the current complete method and the best credible alternative, including self-service, other autonomous providers, conventional software, cooperation, and non-entry where relevant. Investigate what prevents the desired outcome from being delivered economically, reliably, or at the necessary scale.

An Alpha hypothesis specifies the proposed change, its causal mechanism, the outcome that remains protected, its full lifecycle burden, and a decisive failure condition. A lower token bill is not enough if acquisition, verification, integration, delay, or maintenance overwhelms the saving. A technically ordinary component can be part of an extraordinary complete arrangement; novelty is not an exemption from comparison.

The research executive may search backward from the outcome toward missing means, or forward from available capabilities toward newly reachable outcomes. It may transfer a structure from another domain, construct a missing correspondence, or remove the dependency that made the correspondence seem necessary. The choice among those routes is itself an object of investigation, not a fixed strategy imposed on every task.

3.2 Gym: develop the capability before trusting it

A Gym is a mission development environment built from authorized evidence, tools, constraints, and experimental history. It supports proposal, construction, simulation, stress testing, comparison, and correction. Its output is a candidate and a development record, not a certificate of deployment readiness.

The Gym must represent failure modes that matter: missing observations, wrong initial hypothe-

ses, unaffordable tests, incompatible interfaces, misleading references, absent demand, delayed collection, and the possibility that the incumbent is already sufficient. Difficulty should not be manufactured merely to make the new architecture look useful.

Simulation can identify a discriminating physical experiment. It cannot substitute for that experiment when the claim depends on unknown physical behavior. NASA's readiness distinctions between analytical, laboratory, relevant-environment, and operational demonstrations are a useful precedent for keeping those steps separate [8]. We adopt the distinction, not an automatic aerospace certification for this programme.

3.3 Specialist ASI: an earned, bounded designation

In this programme, SPECIALIST ASI is a proposed mission-specific designation. It is not a standardized scientific category or a claim of general superintelligence. An eligible claim must identify the exact candidate, mission distribution, protected test protocol, comparators, resource accounting, practical improvement threshold, safety constraints, uncertainty, and validity conditions. The designation must never travel without that scope.

Freeze the candidate before protected confirmation. Compare it with the current approach and the best credible complete alternative under comparable permissions and resources. Evaluate accepted outcomes, total cost, time, robustness, and material risks. Multiple objectives require declared tradeoffs or non-inferiority conditions; a convenient weighted average must not conceal a disqualifying regression.

No experiment in this paper awards that designation to a model. The cash study evaluates supplied operating arrangements, not AI task performance. The local state machine tests acceptance logic, not the truth of an ASI claim. A label cannot supply the evidence it names.

3.4 SUCCESSOR Ω: preserve, govern, renew

Following sufficient evidence and explicit customer authorization, a capability may become part of the operating institution. The institution preserves legitimate learning, executes within a bounded grant, monitors drift, records consequential changes, and supports expiry, suspension, recovery, and renewal.

These functions form a feedback relationship, not an irreversible staircase. Operational evidence may return a capability to the Gym. A supplier replacement may require only a narrow justified requalification, or a complete rebuild. A changed mission may invalidate the original abstraction. A stopped capability can remain informative without remaining active.

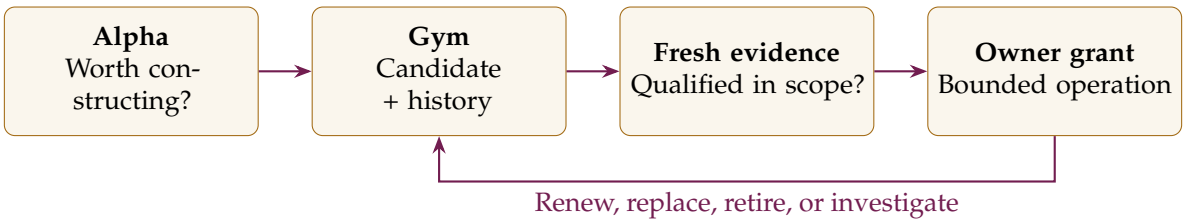


Figure 1: The four-stage vocabulary is implemented through distinct decisions. The development environment does not certify itself; evidence does not issue its own operating permission.

4 Preserve the knowledge; renew the warrant

Let a qualification context be

$$c = (m, s, v, e, q, p), \tag{2}$$

where the entries identify the mission, source material, supplier configuration, environment, evaluator, and policy. A practical implementation must bind the complete executable dependency closure: prompts, adapters, libraries, tool versions, retrieval state, and relevant configuration, not only a short model name. The six strings in our local witness stand for that closure; they do not discover it.

A current evidence record binds a context digest, supported actions, validity interval, test protocol, and outcome. An authority record additionally binds an owner decision, an allowed action set, and its own validity interval. A historical record says what was established under its original conditions. It is not rewritten when current qualification expires.

Definition 1 (Succession without privilege inheritance)**.** *A restoration or successor initialization may import authorized knowledge, executable artifacts, and historical records. It initializes current evidence and active authority to empty. New operation requires an accepted context-bound qualification and a separate explicit grant.*

This conservative default does not prohibit efficient reuse. A validated generic check may remain applicable when its artifact, dependencies, checker, and environmental assumptions are unchanged. Current mission inputs, outputs, and material acceptance requirements still require current checks. A justified equivalence relation can reduce redundant work; a cache key alone cannot establish semantic equivalence.

Proposition 1 (Context-bound action invariant)**.** *Suppose every external action is mediated by a monitor that requires both a passing evidence record and an unexpired explicit grant, binds both to the current context, restricts the grant to the supported evidence scope, and clears the grant whenever qualification is replaced or revoked. If initialization and restoration create neither record, every permitted action is supported and authorized in its current recorded context and interval.*

Proof. Initialization permits no action. Qualification alone creates no grant. A valid authorization transition creates only a subset of the supported actions, in the same context and within the evidence interval. The action check requires those predicates at the time of use. Replacement clears the grant; a context change fails the identity check; expiry or revocation fails the temporal or grant check. Induction over the permitted transitions establishes the invariant. This is a property of the specified monitor, not a proof that its evidence is truthful or that an implementation cannot be bypassed. □

4.1 The assurance system is itself fallible

A hash establishes byte identity under its cryptographic assumptions; it does not establish that the bytes describe the right physical object. Distinct role names do not establish independent organizations, independent model errors, or uncompromised storage. A source can be consistently wrong. A test population can omit the failure that matters. An honest verifier can have an incomplete specification.

The production design therefore needs threat modeling, complete mediation, authenticated owner decisions, isolation of protected evaluations, dependency integrity, recovery tests, and independent validation of evidence semantics. These requirements belong in implementation

qualification. The NIST AI Risk Management Framework provides a relevant governance and lifecycle precedent, without turning our local prototype into certified compliance [10].

The same discipline applies to the Gym. Builder-generated tests can improve development. They cannot become the final judge merely because they favor the builder. Exposed confirmation cases become development material, not a fresh test set for the next renamed candidate. Repeated attempts consume a declared programme-wide error allocation when statistical claims are made.

4.2 Three records, three different questions

A useful implementation keeps a knowledge record, an evidence record, and an authority record conceptually separate. The first answers what the institution can legitimately reuse. The second answers what has been established for this context. The third answers what the owner has actually permitted. A fourth, financial record determines whether the action can be funded without consuming protected obligations.

None is a substitute for another. Scientific correctness cannot create a customer contract. A customer contract cannot establish physical performance. A large balance cannot grant access. A successful delivery cannot retroactively erase an invalid experiment. Institutional continuity consists partly of retaining these distinctions when components and personnel change.

5 Renewal-adjusted Alpha

Let π be a complete construction and operating policy over horizon H. Let Y_t be accepted outcomes and C_t the costs and losses actually incurred or explicitly modeled at time t. For a declared mission utility U and discount factors d_t, define

$$J_H(\pi) = \mathbb{E}\left[\sum_{t=0}^{H-1} d_t\{U(Y_t) - C_t\} - d_H O_H\right], \tag{3}$$

where O_H is the remaining terminal obligation under a stated valuation convention. Monetary costs and nonmonetary utility must be converted only through an explicit owner-approved decision model, or retained as separate constrained objectives.

For a frozen family $\mathcal{B}$ of credible complete alternatives,

$$\alpha_H(\pi) = J_H(\pi) - \sup_{b\in\mathcal{B}} J_H(b). \tag{4}$$

We call this *renewal-adjusted Alpha* when the policies and their costs explicitly include replacement, migration, requalification, downtime, maintenance, and retirement, as well as initial development and operation. The expression is a decision criterion. An estimate of it is not established advantage, and a maximum over a finite supplied family is not a global economic optimum.

5.1 Renewal can reverse a local ranking

Suppose one specialist produces greater contribution on each active period but must be reconstructed whenever its supplier changes. A portable capability produces slightly less contribution but can be requalified at lower cost. A one-period comparison favors the specialist. A complete horizon may favor the portable capability. The decisive object is the supported productive trajectory, not the component's isolated score.

Conversely, portability can be wasteful. Its setup cost may never be recovered. The environment may not change. Retained assets may demand expensive maintenance. A new supplier may be incompatible with the supposed abstraction. An ordinary direct method may dominate. The institution should discover when to build continuity and when not to pay for it.

5.2 Cash is a constraint, not a synonym for profit

Let B_t be cash before a period, u_t its upfront payment, r_t a receipt collected only after accepted delivery, and ρ_t a protected reserve. In the finite study,

$$B_t - u_t \geq \rho_t, \qquad B_{t+1} = B_t - u_t + r_t. \tag{5}$$

The first inequality is checked before collection. The second cannot make an unaffordable action valid retroactively. A positive horizon surplus does not establish that the policy can be started.

The study reports *net cash surplus*, not accounting profit. There is no debt, equity issuance, receivable, tax model, salvage value, or noncash asset valuation. All cash flows are supplied by the fixture. Actual enterprise reporting must additionally reconcile accrual profit, liabilities, taxes, capital expenditures, and unrestricted cash without double-counting them. Agent development consumes actual resources; the absence of a human employee is not the absence of cost. Human labor is charged when it occurs, not inserted as a universal requirement.

5.3 Value capture is a separate construction problem

An owner needs both a useful capability and a voluntary route to receiving enough value to sustain it. Buyers can decline, use another provider, or perform the work themselves. Suppliers can change terms. Competitors can imitate. Owning a portable institution may reduce switching exposure, but its commercial benefit must survive real prices and actual behavior.

The appropriate contract might sell accepted outcomes, continuity of a bounded business function, a qualified instrument, or a separately priced renewal service. These are hypotheses, not measured demand. Customer data and methods do not become the provider's transferable property merely because they could improve another institution. A federation of institutions should exchange only permitted artifacts and evidence, with independent grants on each side.

The durable advantage we seek is the ability to keep constructing worthwhile capabilities, not a permanent monopoly margin. The previous FLYWHEEL Ω counterexample already shows why pricing claims must follow the choices participants can actually compose [1]. Institutional ownership adds another strategic choice: the ability to leave a supplier without abandoning the mission.

6 An exact study of the institution that survives replacement

The study is a deliberately small, fully enumerable construction. Its purpose is to expose a lifecycle mechanism and its failure conditions. We supply the capabilities, contributions, replacement schedule, and action vocabulary. No model discovers these primitives; no customer purchases anything. The experiment tests exact planning, cash feasibility, and the consequences of retaining a useful artifact under those assumptions.

6.1 The frozen environment

There are 16 periods, divided into four supplier epochs of length four. Opening cash is 100 and the protected reserve is 20. The period is an abstract decision interval, not a month or a year. Each period permits one action. Payment precedes any collection. Development, renewal, and retirement consume a period during which no service is delivered.

The institution begins with a legacy service already qualified and available to every arm. That endowment is deliberate: the comparison does not charge the baseline an invented setup burden. A provider-specific specialist can be constructed for 40 units. A portable capability can be constructed for 60 units and incurs one unit of maintenance in every period in which it is retained, including its construction and renewal periods.

Table 1: Supplied nominal action contracts. Maintenance is additional when the portable asset is retained. All values are abstract cash units, not measured prices.

Action	Upfront	Receipt	Availability and effect
Legacy delivery	10	22	Initially available; contribution 12.
Build specialist F	40	0	Qualifies F for the current supplier epoch.
Deliver with F	12	42	Requires current F qualification; contribution 30.
Build portable P	60	0	Retains P and qualifies it for the current epoch.
Deliver with P	10	38	Requires current P qualification; contribution 28 before maintenance.
Renew P	10	0	Retained P receives current-epoch qualification.
Idle	0	0	Maintenance still applies if P is retained.
Retire P	0	0	Removes P and future upkeep; no salvage receipt.

At a supplier change, current specialist and portable qualifications expire. The portable artifact remains available for renewal; the provider-specific specialist must be rebuilt. That difference is an authored property of the two contracts. It is not a measured universal property of portable versus provider-specific AI. Constructing a usable portability layer in a real system is precisely what the proposed industrial experiment must establish.

The complete state is (t, B, o, e_P, e_F): period, cash, whether P is retained, and the epochs for which P and F are qualified. Qualification expires implicitly when an epoch identifier differs from the current epoch. The optimizer can choose the legacy service, build or use either specialist, retain or retire P, renew it, or idle. It is allowed to choose neither new capability.

All modeled service deliveries succeed and collect their specified receipt in the same period when demand is enabled. The default schedule is deterministic and known. There are no material capacity limits beyond one action per period, and no uncertain physical effects. These strong assumptions make the result exactly checkable; they also limit its commercial interpretation.

6.2 Four comparisons, not a protected protagonist

The *full-option planner* has all actions. The *no-portable planner* has the same cash rules and may rebuild F, use legacy, or idle; it tests the value of adding the supplied portable construction. The *legacy-only control* uses the initial service or idles. Finally, an *ordinary equal-content planner* has the same full action set and information. It reaches the same optimum. The ordinary control is not prohibited from retaining knowledge or constructing a portability layer.

We also evaluate a supplied *always-retain policy*: immediately build P, renew it at each epoch change, and otherwise deliver with it. This policy is not allowed to borrow. It returns infeasible when it cannot prefund a required action. Its purpose is to test the tempting but incorrect rule

that institutional continuity should always be purchased.

6.3 Result: the stronger local component loses the horizon

Table 2: Nominal exact result. The 100-unit initial balance is not revenue. The 20-unit reserve remains inside terminal cash.

Policy	Collections	Outflows	Final cash	Cash surplus
Legacy only	352	160	292	192
Best without P	504	304	300	200
Full option set	456	226	330	230
Ordinary equal-content	456	226	330	230

The nominal full-option policy builds P in period zero, delivers three times, then renews and delivers three times in each subsequent epoch. The no-portable optimum builds F at the start of each epoch and delivers three times. P has lower active-period contribution, yet retains 30 more cash units over the complete horizon. Its 230-unit surplus is 15% higher than the 200-unit no-portable surplus, *within this fixture*. This is not a measured business margin or an AI productivity estimate.

The result has a useful second interpretation. The winning route collects *less* revenue: 456 rather than 504. It wins because it spends less on reconstructing the means of delivery. Maximizing transaction value, isolated performance, or gross collections would not by itself identify the better funded trajectory.

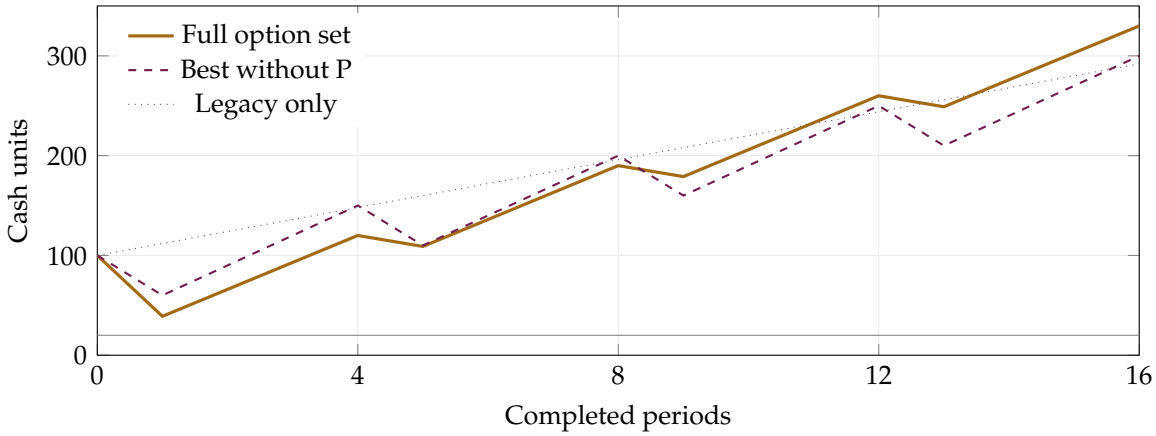


Figure 2: The exact cash paths, including construction and renewal interruptions. No future collection pays for its own upfront cost. The horizontal line marks the protected reserve.

6.4 The crossover is explicit

For a fixed sequence of n complete epochs of length L, suppose rebuilding costs D, portable construction costs B, renewal costs R, and upkeep is m per period. Let g_F and g_P be active-period contributions before portable upkeep. Each construction or renewal consumes one period. Conditional on financing both sequences,

$$S_F(n) = n[(L-1)g_F - D], \tag{6}$$

$$S_P(n) = n(L-1)g_P - B - (n-1)R - nLm. \tag{7}$$

Hence

$$S_P - S_F = n\underbrace{[(L-1)(g_P - g_F) + D - R - Lm]}_{\text{per-epoch continuity benefit}} - B + R. \tag{8}$$

This is a comparison of two specified sequences, not proof that either is optimal in every environment. The exact planner establishes the nominal optimality separately.

With the nominal parameters, $S_P - S_F = 20n - 50$. P loses by 30 in one epoch, loses by 10 in two, wins by 10 in three, and wins by 30 in four. The advantage is created by repeated avoided reconstruction. It is not a mysterious multiplier of intelligence.

A distinction worth remembering

The best component is not necessarily the best institution. The best institution is not necessarily the most persistent one. Preserve a capability only when its supported future uses justify construction, renewal, and upkeep.

6.5 The sensitivity study includes losses and non-entry

Before executing the grid, the protocol fixed horizons $H \in \{8, 12, 16\}$, epoch lengths $L \in \{4, 8\}$, opening cash $C \in \{40, 70, 100\}$, renewal costs $R \in \{5, 10, 25\}$, and upkeep $m \in \{0, 1, 3\}$. Their Cartesian product contains 162 environments. Both optimizers were checked against independently implemented forward reachability in every environment: 324 exact comparisons.

Adding the portable option strictly improves terminal cash in 19 environments and ties in 143. The largest gain is 61 units. There are no negative optimum differences because the full feasible set contains the no-portable set. *That nonnegativity is a mathematical inclusion property, not evidence of robust positive treatment effects in a population of businesses.* The grid is supplied, not sampled from a real market.

Forced retention performs much worse: it is infeasible in 108 environments, loses to the no-portable optimum in 43, wins in 11, and ties in none. Infeasibility is an outcome, not a missing observation. These counts show why the institution must retain the option to delay, simplify, retire, or decline.

Table 3: Diagnostic cases, including explanatory additions. Values are terminal cash; opening cash changes where stated. The full planner is free not to construct P.

Changed condition	Full	Without P
No supplier turnover during 16 periods	510	510
No demand and no obligation to build	100	100
Portable upkeep rises to 5	300	300
Portable renewal rises to 40	300	300
Opening cash is only the reserve, 20	20	20
Opening cash is 40	238	238
Horizon is one four-period epoch	150	150

With no turnover, the higher-contribution F route wins. With opening cash equal to the reserve, even a profitable legacy delivery cannot be prefunded. With a four-period horizon, the always-portable sequence ends at 120 while the optimum ends at 150. The favorable nominal result does not license an unconditional preference for portability.

6.6 Independent computation and its limits

The primary solver uses memoized backward induction. The second implementation enumerates all reachable states forward and re-specifies action effects independently, without calling the primary transition function. The two implementations agree on every registered grid comparison. They share the authored specification and constants; agreement cannot validate those economic assumptions.

The companion also contains 35 unit tests covering planning and a local qualification/authority state machine. The latter preserves knowledge across export and restoration while rejecting inherited active proof or authority. Tests exercise changed context fields, expiry, revocation, unsupported actions, and missing owner approval. They establish behavior of this local software, not a hardened permission service or independent certification.

The implementation limit of 32 periods is a declared finite test boundary, not a business-scale ceiling. No live frontier-model calls, real customer data, actual collections, or physical operations occur. CPU time spent on publication and software construction is not hidden inside the abstract model's cash values. The supplied build, renewal, and maintenance prices are model inputs, not measured complete costs of this research programme.

7 What must remain renewable

Useful capability is not merely installed capability. It must remain technically applicable, adequately evidenced, permitted, and affordable. This creates a renewal problem even when the underlying code remains unchanged.

7.1 Renewal is a capacity constraint

Suppose each deployed unit requires a current certificate, each certificate supports at most τ units of time, and the verification system can issue at most v such unit-certificates per unit time. Let B_0 denote initially available remaining certified unit-time. For a horizon of length T, authorized productive unit-time is bounded by

$$\int_0^T N_{\text{active}}(t)\,dt \leq B_0 + \tau v T. \tag{9}$$

Thus long-run average authorized deployment cannot exceed τv under these assumptions. Adding unqualified machines does not remove that bound.

This is a counting bound for a declared certification policy, not a physical law. A valid shared proof, a better production process, or a different independently justified validation architecture can change the unit of certification, τ, or v. The productive invention may be the construction that makes those changes legitimate. It is not the deletion of the requirement while pretending reliability has been preserved.

7.2 Invent the renewal mechanism

A recurrent qualification burden can be approached in several ways. Improve the test. Construct an interface that makes irrelevant changes unable to affect the protected property. Preserve evidence during production so it need not be reconstructed afterward. Replace a fragile dependency. Or determine that full replay remains cheaper than maintaining a specialized reuse mechanism.

These routes transform the conditions under which intelligence becomes useful. The right target is not the minimum number of checks. It is the lowest complete burden of maintaining the required accepted outcome, including missed failures and the cost of invalidating stale conclusions.

A proof-producing interface may have large enabling value because many successors can use it. That does not make its value infinite. Integration, dependency tracking, compatibility testing, and obsolescence consume resources. The institution must compare its continued use with a clean rebuild and with simpler direct methods.

7.3 A renewal reserve is not a growth forecast

The operational reserve in the toy model is a supplied constant. A real institution needs a justified policy for existing obligations, rollback, incident response, and expected renewal work. Protecting those funds can prevent reinvestment even when accounting profit is positive. That is a cash decision, not a reason to erase the profit result.

Likewise, available cash does not establish an attractive next project. The proposed investment must identify what will be built, what independent acceptance would establish, when cash will be consumed, which liabilities remain, and why the complete route is preferable to alternatives. If these conditions fail, preserving optionality can be the better action.

8 The institution must improve more than its inventory

The numerical study establishes the value of a supplied retention option in particular environments. It does not establish that a machine has become better at inventing that option. This boundary is central to the next research step.

Prior research supplies substantive but bounded precedents. DreamCoder jointly develops reusable program abstractions and search machinery in evaluated synthesis domains [5]. Enhanced POET studies the generation of challenges and solutions with transfer between environments [12]. Automated Design of Agentic Systems constructs executable agent designs rather than selecting only a fixed workflow catalogue [6]. AlphaEvolve uses generated code and automated evaluation for scientific and algorithmic discovery [11]. The Darwin Godel Machine explores an archive of self-modifying coding agents [13]. These works motivate component mechanisms; none is evidence for the complete customer-owned, economically sustained, physically expanding institution proposed here.

8.1 Three levels of inheritance

A retained *artifact* gives a successor a tool. A retained *applicability model* helps determine where the tool is appropriate. A retained *constructor* generates related tools under new conditions. The transition between each level is a separate empirical claim.

A failed transfer can be as informative as a success. It may show that the analogy was invalid, identify the missing interface, or reveal that the original objective does not actually need the assumed dependency. Negative evidence must retain its scope. Failure within one budget is not proof of universal impossibility; a success in one mission is not universal qualification.

The most ambitious inheritance is a better procedure for discovering constructors. It must be evaluated as a discovery policy, not inferred from the quality of one descendant. Replacing a

costly committee with one good discriminating experiment can be an improvement. Architectural complexity is not the objective.

8.2 The crossed experiment

Let η_0, η_1 be earlier and revised discovery procedures, and K_0, K_1 earlier and updated starting capabilities. Evaluate all four combinations on genuinely fresh mission campaigns:

	Earlier capabilities K_0	Updated capabilities K_1
Earlier procedure η_0	Reference	Inherited-equipment benefit
Revised procedure η_1	Discovery-policy benefit	Combined effect

Each arm receives comparable permissions and complete budgets. Strong unrestricted conventional construction and ordinary equal-content reuse remain available controls. Count the historical cost of acquiring the proposed improvement as well as its new campaign cost. The comparison unit for a claim about discovery is the complete campaign, not thousands of correlated substeps presented as independent replications.

An intervention is promoted only for the scope and claim that the evidence supports. A useful artifact can be adopted without proving improved discovery. A profitable business can expand without improving its intelligence. Conversely, a valuable scientific instrument may deserve a bounded research investment before it generates revenue. Keeping these cases separate allows ambition without a fictional universal success gate.

8.3 The research executive can outgrow this paper

The four-stage vocabulary is a public map, not a permanent grammar of intelligence. An effective institution may find a new representation, a different sequence of experiments, a simpler commercial arrangement, or a better way to allocate search. Its own route to improvement should remain revisable under the same protected purpose and evidence rules.

A genuine advance changes what becomes discoverable, constructible, verifiable, or economically deliverable. A new name for a fixed workflow does not. The complete mechanism must be executable and its contribution must survive a comparison in which the alternative is allowed to be competent.

9 From productive surplus to a physical frontier

An institution of intelligence becomes consequential at civilization scale only if its work changes what people can actually produce, understand, and sustain. Financial wealth is one means of coordination. It is not energy, a functioning factory, a validated material, or a beneficial outcome.

The proposed path has distinct links:

$$\text{accepted value} \longrightarrow \text{available earned resources}$$
$$\longrightarrow \text{qualified construction} \longrightarrow \text{commissioned capacity.}$$

Knowledge can strengthen every link, but cannot make the links identical. The earlier production-frontier paper explicitly distinguishes these quantities [2]. SUCCESSOR Ω adds the requirement that the institution responsible for them remain capable of renewal across component generations.

9.1 An auditable conversion contract

For an investment x, define a proposed capacity-conversion contract

$$\mathcal{C}_x = (Y_x, P_x, F_x, Q_x, T_x, R_x), \tag{10}$$

where Y_x is the intended physical output; P_x the production and dependency plan; F_x the timed funding and obligation schedule; Q_x the independent qualification route; T_x construction, integration, and deployment timing; and R_x repair, renewal, retirement, and residual obligations. A project becomes commissioned capacity only when the relevant output and acceptance have actually occurred.

This contract is not a new financial instrument or a claim of available counterparties. It is a research and operating record. It prevents a forecast of future inventions from entering the ledger as installed machinery, and prevents a purchase commitment from entering it as available cash. Failed projects retain their costs even when the portfolio changes direction.

9.2 The frontier is constrained by complementary flows

If each commissioned unit requires material a, manufacturing work b, test capacity c, and deployable funding d, then a simplified steady-flow bound is

$$\dot{N} \leq \min\left\{\frac{\dot{M}}{a}, \frac{\dot{W}}{b}, \frac{\dot{Q}}{c}, \frac{\dot{F}}{d}\right\}. \tag{11}$$

Units and time horizons must be compatible. The equation omits inventories, delays, reliability, resource substitution, and network topology; a detailed design must restore them. Its purpose is to expose the error of extrapolating one growing resource while ignoring the complementary requirements.

An intelligence can attempt to change a denominator, construct another route, increase a limiting throughput, or design a process that eliminates a requirement. It cannot assume that the change has occurred. Financial reinvestment with no material-processing capacity may accumulate claims rather than productive machinery. Manufacturing with inadequate acceptance capacity may accumulate uncommissioned inventory.

An expanding institution should therefore select inventions partly by which consequential bottleneck they remove. It should also ask whether the bottleneck moves elsewhere after the change. The best next invention can be a measurement, a production interface, or a coordination arrangement rather than the most visibly sophisticated component.

9.3 A stellar-scale horizon, stated in physical units

The distant research horizon is a productive civilization able to use a substantial fraction of stellar energy while preserving safety, legitimacy, and human benefit. This is an ambition to organize research, not a capability delivered by this paper or a date predicted by its model.

The IAU nominal solar luminosity is $L_\odot^{\mathrm{N}} = 3.828 \times 10^{26}\,\mathrm{W}$; its nominal solar irradiance is $S_\odot^{\mathrm{N}} = 1361\,\mathrm{W\,m^{-2}}$ [7]. These are nominal conversion constants, not claims about the Sun's exact instantaneous output. With intercepted fraction f, conversion efficiency η, and those reference values,

$$P_{\mathrm{use}} = \eta f L_\odot^{\mathrm{N}}, \qquad A_{\mathrm{collector}} \geq \frac{f L_\odot^{\mathrm{N}}}{S_\odot^{\mathrm{N}}}. \tag{12}$$

The area is a simplified effective intercepting area at the reference irradiance. It assumes ideal orientation and no overlap or shading. It is not the surface area of a realizable structure.

For illustration only, take $f = 10^{-12}$ and $\eta = 0.30$. The intercepted power is $3.828 \times 10^{14}\,\mathrm{W}$, useful power is $1.1484 \times 10^{14}\,\mathrm{W}$, and effective collecting area is approximately $2.81 \times 10^{11}\,\mathrm{m}^2$, or 281,264 square kilometres. An assumed areal mass of $1\,\mathrm{kg\,m^{-2}}$ would already imply approximately $2.81 \times 10^{11}\,\mathrm{kg}$ of collector material alone. Neither efficiency nor areal mass is a demonstrated design in this programme.

Heat rejection creates another explicit requirement. Under a deliberately simplified cold-background radiator model,

$$A_{\mathrm{radiating}} \geq \frac{P_{\mathrm{waste}}}{\epsilon \sigma T^4}, \tag{13}$$

where $\sigma \simeq 5.670374419 \times 10^{-8}\,\mathrm{W\,m^{-2}\,K^{-4}}$ [9]. Taking $T = 300\,\mathrm{K}$, emissivity $\epsilon = 0.9$, and $P_{\mathrm{waste}} = 0.7 f L_{\odot}^{\mathrm{N}}$ gives approximately $6.48 \times 10^{11}\,\mathrm{m}^2$ of emitting area for the same illustrative fraction. This is total radiating area, not automatically the footprint of a one- or two-sided panel. Background radiation, solar heating, view factors, transport, degradation, and the eventual dissipation of useful energy are not modeled.

Why keep the physical arithmetic?

Even a trillionth of a star's nominal luminosity implies a vast material, thermal, manufacturing, and coordination problem. A forecast of compounding money is not a construction plan for it. The opportunity is to invent the successive means that make each genuine step buildable.

9.4 Human benefit is not measured in watts alone

Greater capacity can improve health, scientific access, resilience, and material well-being. It can also create concentrated power, ecological damage, or catastrophic exposure. The mission must therefore preserve explicit constraints on harm, distribution, consent, and accountability instead of assuming that greater throughput is inherently beneficial.

A customer-owned institutional design should make control and exit practical rather than concentrating authority in an irreplaceable supplier. At larger scales, coordination among independent institutions must remain permissioned, contestable, and subject to applicable public governance. The architecture is not a proposal for one autonomous owner to acquire unbounded authority. Greater capability increases the importance of justified limits and recoverable operation.

10 The next consequential programme: an Industrial Requalification Foundry

The proposed flagship is a foundry that constructs *portable, independently testable industrial capabilities* across changing instruments, model providers, and operating conditions. The intended outcome is not another evaluation dashboard. It is reduced lifecycle burden for useful scientific and industrial work that currently loses too much value at each transition.

An initial candidate mission is evidence-preserving experimental planning for battery-process qualification. The institution would connect source-grounded materials knowledge, simulation, experimental design, measured manufacturing variation, acceptance tests, and renewal after an

instrument or supplier change. This is a candidate, not a market selection already established by data. Other missions may offer greater consequence or a sharper initial test.

10.1 Construct a missing means, not a narrative

One prospective instrument is a *renewal-aware experiment compiler*. It would map a proposed configuration change to the claims potentially affected, preserve justified unchanged evidence, construct the smallest discriminating experiments needed for new claims, and produce an executable qualification plan with complete costs. The plan can conclude that the dependency closure is unknown or that a clean rebuild is preferable.

The critical inventive step is constructing the mapping from changes to consequential claims, not merely sorting a supplied dependency graph. A prototype that receives every causal edge from its designer should be reported as a supplied-graph implementation. An autonomous discovery claim requires the system to infer or construct missing structure and test it against protected cases.

A second prospective instrument would make work easier to qualify by construction: typed experimental interfaces, provenance-preserving observations, declared calibration contracts, and explicit failure behavior. That could change the amount of recurrent reasoning required. It also creates costs and new failure modes, which belong in the comparison.

10.2 The first decisive campaign

Start with one bounded mission and three materially different routes: a strong direct specialist; ordinary modular engineering with full documentation and reuse; and a candidate permitted to invent a new representation, experiment, or interface. All receive the same substantive source access, construction permissions, and complete campaign budget. The candidate must not receive an exclusive prebuilt connector.

Separate development cases from fresh confirmation. Include cases where the change is irrelevant, where a narrow test is sufficient, where complete replay is necessary, where a reference is coherently wrong, and where no affordable test justifies operation. Supplier replacement, source mutation, changed measurement semantics, and expired permission are different interventions and must not be collapsed into one generic drift score.

The primary endpoint is the cost and time to an independently accepted mission outcome, subject to fixed material quality and safety requirements. Report wrong approvals, false rejections, uncertainty, interruption recovery, and actual consumed resources. For a physical claim, a simulation result advances only to a justified physical test; it does not count as that test. The reserve for independent confirmation is protected before search begins.

Then test a fresh successor, such as another production line or measurement configuration. Compare inherited machinery with no inheritance and ordinary equal-content reuse. Finally, perform the crossed discovery-policy comparison from Section 8. These are three separate experiments, not three names for the same favorable run.

10.3 The voluntary commercial bridge

The buyer hypothesis is that a qualified institution can deliver a valuable bounded outcome and preserve the means of delivering it through future changes. Buyer acceptance must establish actual willingness to pay for that outcome relative to alternatives, not agreement with the vision. Proposed prices are experiments, not evidence of demand.

A pilot should record authorized onboarding, independently accepted work, actual collection, acquisition burden including declined offers, failures, renewal obligations, and unrestricted funds. Repeat purchase and portability after supplier replacement are separate observable outcomes. Commercial performance can be recognized without claiming improved invention; a technically impressive renewal method can be rejected when a simpler service produces greater complete value.

Physical expansion follows only through separately funded and qualified projects. A positive digital-service surplus may finance laboratory work. It does not establish that the laboratory work will succeed. An industrial result may enable another capability. That enabling relationship requires its own comparison, and the full lineage retains unsuccessful costs.

11 Scope, failure modes, and what would falsify the thesis

Our strongest established numerical claim is conditional: for the registered finite environments, the exact solver and the independent reachability implementation agree, and the nominal portable route has greater cash surplus than the best supplied no-portable route. The ordinary equal-content planner ties it. The state-machine tests establish local transition behavior. The physical arithmetic is a dimensional scale check. None is evidence of autonomous commercial success or a newly achieved level of intelligence.

The economic fixture gives P cheaper renewal by assumption. It does not estimate how often suppliers change, how well capabilities transfer, what buyers will pay, or the probability of qualification failure. Known future turnover is especially favorable to planning. A real implementation must confront uncertain change, collection delay, competing prices, obsolescence, multi-period experiments, and operational mistakes. The study's positive cases do not establish their prevalence.

The institutional hypothesis loses its practical value in a mission where the strongest ordinary arrangement already preserves equivalent capability at equal or lower complete cost. It also fails when export omits a decisive dependency, renewal produces stale assurance, maintenance absorbs the operating surplus, the owner cannot actually control the assets, or preserved machinery does not help fresh work. Those are substantive negative results, not reasons to weaken the comparator.

The broader recursive hypothesis requires more. A history of successful tools may still leave discovery unchanged. A revised search procedure can look better because it consumes more resources or sees confirmation material. A larger archive can make retrieval worse. The design must allow a simple conventional alternative to win and the incumbent to remain in service.

Finally, no general inference runs from intelligence to unlimited growth. Production needs physical resources, economic participation, infrastructure, time, and legitimate authority. A useful method can accelerate a link without establishing the entire chain. The programme should pursue the links that could have the greatest consequence, while reporting exactly which have been built and which remain hypotheses.

12 Conclusion: build what can build beyond itself

A successful answer is an output. A useful invention is a capability. An enduring institution preserves the legitimate means of producing further capabilities while remaining able to replace the components that once made it successful.

SUCCESSOR Ω puts that distinction into an economic and experimental form. Find an advantage worth constructing. Develop it in a mission environment. Make the claim survive fresh comparison. Grant bounded authority separately. Preserve what genuinely transfers, renew what must be current, and retire what no longer earns its cost.

The exact study demonstrates why the lifecycle matters: a lower-contribution portable capability can beat repeated reconstruction, yet continuity can also be unaffordable or inferior. Its ordinary equal-content tie identifies the actual mechanism without inventing superiority for a label. The research challenge is now to discover, build, and qualify such enabling arrangements in consequential missions.

The larger opportunity is to compound more than money or artifacts. It is to develop institutions whose work expands what can be investigated, constructed, delivered, and sustained—and whose earned resources help make the next genuine advance possible. Scientific instruments, industrial systems, and energy-producing capacity are destinations for that process, not numbers to be asserted in advance.

The governing ambition

Own the institution. Replace the models. Preserve the mission.

Make the next breakthrough easier to construct—without making its proof easier to evade.

A Formal details and boundary conditions

A.1 Finite optimization

For state $s_t = (t, B_t, o_t, e_{P,t}, e_{F,t})$, let $\mathcal{A}(s_t)$ be the legal supplied actions satisfying the pre-collection reserve. Each action has deterministic cost $u(s_t, a)$, collection $r(s_t, a)$, and successor state $T(s_t, a)$. The primary recursion is

$$V_H(s_H) = B_H, \qquad V_t(s_t) = \max_{a \in \mathcal{A}(s_t)} V_{t+1}(T(s_t, a)). \tag{14}$$

The implementation breaks equal terminal-cash outcomes by lexicographic action-sequence order. That affects the selected trace, not the optimal cash value. Idle is feasible when no portable asset is retained. When upkeep is unaffordable, retirement remains feasible and consumes the period. Thus an admitted state always has a continuation.

Proposition 2 (Exactness within the declared action set)**.** *For the finite deterministic model, backward induction returns the maximum terminal cash over all reserve-respecting action sequences. Independent forward enumeration returns the same maximum when it uses the same transition specification.*

Proof. At the terminal layer the value equals cash by definition. If continuation values at $t+1$ are exact, maximizing over every legal first action at t examines every feasible sequence partitioned by that action. Induction proves the first statement. Forward enumeration starts from the admitted state and adds every legal successor at each period; induction identifies its layer with exactly the reachable state set. Taking maximum cash at the final layer proves the second. Agreement between implementations is a useful defect check, not evidence that the supplied model describes a market. □

Corollary 1 (Option inclusion)**.** *If $\mathcal{A}_0$ is a subset of $\mathcal{A}_1$ at every state and all common action effects, endowments, and constraints agree, the optimum under $\mathcal{A}_1$ cannot be lower.*

Proof. Every feasible $\mathcal{A}_0$ action sequence is also feasible under $\mathcal{A}_1$. Maximization over the larger set contains the original maximum. □

This explains the absence of negative full-option differences. It does not protect the always-retain policy, which removes choices and pays costs even when they are not justified. Nor does it justify treating a preconstructed portability layer as free in a comparison of discovery systems.

A.2 The renewal crossover

If every epoch contains one construction or renewal period and $L-1$ active periods, F earns $(L-1)g_F - D$ per epoch. P earns $n(L-1)g_P$, pays its initial construction B, pays $n-1$ renewals, and pays nL periods of upkeep. Subtraction gives Equation 8.

When $\Delta = (L-1)(g_P - g_F) + D - R - Lm > 0$ and $B > R$, P strictly wins those fixed sequences when

$$n > \frac{B-R}{\Delta}. \tag{15}$$

The first strictly winning integer is $\lfloor (B-R)/\Delta \rfloor + 1$. If $\Delta \leq 0$, additional epochs do not create the same amortization advantage. Other signs of $B - R$ must be evaluated directly from the inequality, rather than forcing this positive-threshold interpretation.

Let D_P^{extra} and D_F^{extra} be additional total lifecycle burdens omitted from the toy contracts. Under a compatible monetary accounting convention, the nominal 30-unit difference survives only

if $D_P^{\text{extra}} - D_F^{\text{extra}} < 30$. Liquidity can fail even before that surplus threshold: the timing of an additional burden matters as well as its total. This sensitivity statement does not estimate those burdens.

A.3 Prefix financing

For an ordered cash sequence with constant reserve ρ, define the cumulative pre-collection deficit through action j as

$$D_j = \sum_{i=0}^{j} u_i - \sum_{i=0}^{j-1} r_i. \tag{16}$$

The sequence is feasible exactly when its initial cash satisfies

$$B_0 \geq \rho + \max(0, \max_j D_j). \tag{17}$$

Proof. Immediately before receipt r_j, cash equals $B_0 - D_j$. Requiring this quantity to be at least ρ for every j is equivalent to the bound. The extra zero ensures the initial reserve itself is funded. □

For the nominal immediate P construction, the first payment is 61 including upkeep. Opening cash must therefore be at least 81 before this fixed route can begin. A different policy may first earn legacy receipts, delay construction, or choose F. Equation 17 is a route-specific condition, not a universal minimum capital for the mission.

A.4 Renewal-throughput bound

Let certificates issued within $[0, T]$ each contribute at most τ authorized unit-time, and let at most vT certificates be issued. Initial residual coverage is at most B_0. The union of covered productive unit-time cannot exceed the sum of these coverage budgets, yielding Equation 9. Overlap only makes the bound looser. Dividing by T and taking a long horizon gives average active units at most τv if $B_0/T \to 0$.

A certificate covering a fleet is not a unit-certificate under this assumption. A validated fleet-level qualification changes the model and may improve the bound; an unsupported fleet-level assertion does not. Hardware construction throughput, material availability, and repair constraints can impose additional tighter bounds.

B The numerical record

B.1 Complete nominal route

Period	Action	Opening	Paid	Collected	Closing
0	Build P	100	61	0	39
1	Deliver P	39	11	38	66
2	Deliver P	66	11	38	93
3	Deliver P	93	11	38	120
4	Renew P	120	11	0	109
5	Deliver P	109	11	38	136
6	Deliver P	136	11	38	163
7	Deliver P	163	11	38	190
8	Renew P	190	11	0	179
9	Deliver P	179	11	38	206
10	Deliver P	206	11	38	233
11	Deliver P	233	11	38	260
12	Renew P	260	11	0	249
13	Deliver P	249	11	38	276
14	Deliver P	276	11	38	303
15	Deliver P	303	11	38	330

The trace above is generated from the companion rather than copied from a plotted image. The publication build checks the terminal balance, collections, outflows, reserve, grid counts, and physical arithmetic before typesetting.

B.2 Protocol history and stopping rule

The finite economic protocol fixed the primary fixture, comparator permissions, 162-case grid, and qualitative adverse cases before numerical evaluation in this authoring session. It was an internal prospective record, not a timestamped independent preregistration. The one-epoch and 40-unit opening-cash diagnostic traces were added as explanatory checks; the latter setting is already in the frozen grid. No selection claim is made from those additions. A planned additional-burden check is reported algebraically in Appendix A, rather than represented as a separately executed market experiment.

The stopping rule is completion of the registered grid, diagnostics, and regression suite. There is no search for a favorable random seed, no independent sample of businesses, and no confidence interval over an invented population. The final source is bound to its generated results and rerun during publication qualification. Source changes are visible in normal version control; the favourable nominal case is not relabeled as universal performance.

B.3 Cost scope and interpretation

The finite model includes exactly the listed payments, deterministic accepted collections, upkeep, and one-period construction and renewal delays. It excludes real customer acquisition, actual model-provider invoices, uncertainty in demand and qualification, tax and liability accounting, physical deployment, and the cost of creating this paper. Those exclusions are not estimates of zero. They prevent the experiment's cash surplus from being described as demonstrated complete-cost commercial profit.

The nominal initial 100 units are a fixture, not a mandatory dollar capitalization for AGI ALPHA.

The programme can compare different initial financing, staged construction, cooperation, and non-entry, subject to actual authority and voluntary participation. The separate US$3 billion trailing-365-day complete-cost profit objective is not achieved by this paper. Multiplying abstract units by a currency scale would not achieve it either.

C A claim ledger for the complete programme

Claim	Evidence in this edition	What it does not establish
Lifecycle reversal	Exact nominal policies, cash traces, and crossover identity.	Measured demand, actual prices, or autonomous invention.
Finite optimality	Backward solver and independent forward enumeration agree on 324 grid comparisons.	Global optimality over unspecified strategies or markets.
Optional retention	19 positive differences and 143 ties; forced retention often loses or cannot start.	A frequency estimate for real-world portability benefits.
Scoped action logic	Local tests of context, grant scope, expiry, revocation, and clean restoration.	Secure complete mediation, authentic identities, truthful evidence, or production certification.
Institutional ownership	Export/restore semantics and proposed acceptance contract.	Legal rights to third-party assets or a deployed customer product.
Inventive succession	Fresh-successor and crossed-policy experimental designs.	Demonstrated improvement of a discovery procedure.
Commercial compounding	Explicit collection, obligations, and reinvestment requirements.	Actual customer profit or the programme's monetary milestone.
Physical expansion	Conversion-contract framework, complementary-flow bounds, dimensional calculations.	A tested collector, radiator, factory, or industrial growth forecast.
Stellar-scale capability	A long-horizon research objective with physical quantities exposed.	An achieved civilization level, an investment return, or a timetable.

D Reproduce, inspect, extend

The companion requires Python 3.10 or later and the standard library. From the source directory:

```
python3 -m unittest -v test_companion
python3 companion.py --output results.json
python3 build.py --output /tmp/successor-publication
```

The last command also needs a PDFLaTeX/BibTeX installation with the packages named in `paper.tex`, Poppler, and CairoSVG for the share image. It generates the cash data and trace table, compiles the paper, records tool versions and file hashes, and packages a deterministic source archive. The build uses a fixed source-date epoch and suppresses variable PDF metadata. Byte identity is checked in the recorded build environment; equivalent output across arbitrary TeX versions is not assumed.

The archive contains the manuscript, bibliography, public reader template, original vector share artwork, exact model, tests, protocol, and build tooling. It contains no private conversation

documents, customer data, credentials, or copied model weights. The versioned GitHub edition provides the released PDF, source ZIP, and checksums. A source/PDF manifest binds the artifact to the source files that generated it.

To extend the study, create a new declared protocol rather than silently changing this edition's assumptions. Useful next variations include uncertain replacement time, qualification failure, delayed collection, multiple complementary capabilities, finite maintenance capacity, and actual measured migration costs. Any comparison claiming improved discovery must give conventional alternatives the same information and construction permissions.

Authorship and research assistance

Vincent Boucher specifies the research direction and is the named author of this edition. AI-assisted synthesis, mathematical drafting, implementation, and publication tooling were used in its preparation. The finite model and its constants are authored inputs. Independent implementation of the reachability check is an algorithmic cross-check within this publication workflow, not independent external peer review or physical validation.

References

[1] Vincent Boucher. FLYWHEEL ω: The economy as an output. MONTREAL.AI / QUEBEC.AI, November 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/releases/tag/flywheel-omega-v1.0.0. Research edition 1.0.0; 27 pages.

[2] Vincent Boucher. Beyond the production frontier: Transformational intelligence from verified discovery to stellar-scale civilization. MONTREAL.AI research manuscript, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/247d0dbebb6b04bee52bf18bbfe11a9d9a2737c5/docs/research/beyond-the-production-frontier-20261019.pdf. Preserved public edition; 26 pages.

[3] Vincent Boucher. SUPER ALPHA: Inventing economic possibility. MONTREAL.AI / QUEBEC.AI, November 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/releases/tag/super-alpha-v1.1.0. Research edition 1.1.0; 28 pages.

[4] Vincent Boucher. Transformational intelligence: Learning to change the game. MONTREAL.AI research manuscript, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/247d0dbebb6b04bee52bf18bbfe11a9d9a2737c5/docs/research/transformational-intelligence-202610.pdf. Flagship paper; preserved public edition, 24 pages.

[5] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sable-Meyer, Luc Cary, Lucas Morales, Luke Hewitt, Armando Solar-Lezama, and Joshua B. Tenenbaum. DreamCoder: Bootstrapping inductive program synthesis with wake-sleep library learning. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, pages 835–850, 2021. doi: 10.1145/3453483.3454080. URL https://arxiv.org/abs/2006.08381.

[6] Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. arXiv:2408.08435, 2024. URL https://arxiv.org/abs/2408.08435v2. Version 2, March 2025.

[7] International Astronomical Union. Resolution B3: Recommended nominal conversion constants for selected solar and planetary properties. Resolutions adopted at the XXIX General Assembly, 2015. URL https://iauarchive.eso.org/static/resolutions/IAU2015_English.pdf.

[8] NASA. Technology readiness levels. Space Communications and Navigation Program, 2023. URL https://www.nasa.gov/directorates/somd/space-communications-navigation-program/technology-readiness-levels/. Accessed November 2026.

[9] National Institute of Standards and Technology. CODATA value: Stefan–boltzmann constant. The NIST Reference on Constants, Units, and Uncertainty. URL https://physics.nist.gov/cgi-bin/cuu/Value?sigma. Accessed November 2026.

[10] National Institute of Standards and Technology. Artificial intelligence risk management framework (AI RMF 1.0). Technical Report NIST AI 100-1, NIST, 2023. URL https://doi.org/10.6028/NIST.AI.100-1.

[11] Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco J. R. Ruiz, Abbas Mehrabian, M. Pawan Kumar, Abigail See, Swarat Chaudhuri, George Holland, Alex Davies, Sebastian Nowozin, Pushmeet Kohli, and Matej Balog. AlphaEvolve: A coding agent for scientific and algorithmic discovery. arXiv:2506.13131, 2025. URL https://arxiv.org/abs/2506.13131v1.

[12] Rui Wang, Joel Lehman, Aditya Rawal, Jiale Zhi, Yulun Li, Jeffrey Clune, and Kenneth Stanley. Enhanced POET: Open-ended reinforcement learning through unbounded invention of learning challenges and their solutions. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9940–9951, 2020. URL https://proceedings.mlr.press/v119/wang20l.html.

[13] Jenny Zhang, Shengran Hu, Cong Lu, Robert Tjarko Lange, and Jeff Clune. Darwin gödel machine: Open-ended evolution of self-improving agents. arXiv:2505.22954, 2025. URL https://arxiv.org/abs/2505.22954.