Beyond the Production Frontier

Transformational Intelligence

from Verified Discovery to Stellar-Scale Civilization

Vincent Boucher
MONTREAL.AI
Montréal, Québec, Canada

October 19, 2026 · Research manuscript

Abstract

The economically decisive output of an advanced AI may not be a better answer or a greater volume of automated work. It may be a transformation of the system that turns discoveries into further productive capacity. We develop this possibility as a research program for *transformational intelligence*: identifying unnecessarily restrictive problem formulations, constructing different ones, and demonstrating that the change expands what can be achieved. The object of invention is the production network itself—its representations, experiments, process requirements, investment sequence, and means of building the next capability. We distinguish validated knowledge, commercial surplus, installed capacity, and human benefit; derive a resource-cut obstruction and a conditional expansion certificate; and execute a finite capital-formation study. With one seed module, improving either of two complementary production requirements alone cannot initiate expansion. Changing both, paying their costs, and observing a qualification delay yields nineteen installed modules after six periods while preserving all six required service deliveries. A separately implemented enumerator confirms the optimum. A second reformulation changes investment timing: a three-unit initial treasury that cannot buy the joint improvement immediately reaches ten modules by waiting for realized receipts. These are exact results in supplied synthetic models, not AI-discovered technologies or measured industrial growth. We then connect the mechanism to a falsifiable path from valuable digital artifacts to energy, manufacturing, and stellar-scale infrastructure, with explicit luminosity, mass, heat-rejection, and lead-time accounting. AGI Alpha Agent and SUCCESSOR Ω provide the named implementation target, not evidence that the distant endpoint has been achieved. The central thesis is that intelligence can be invested in changing the conditions of future achievement—but every conversion from an idea to wealth, and from wealth to physical capacity, must earn its own evidence.

Keywords: transformational intelligence; productive capacity; production networks; endogenous growth; capability construction; SUCCESSOR Ω.

Find your Alpha. Build your SUCCESSOR Ω.
github.com/MontrealAI/AGI-Alpha-Agent-v0

1 The growth process is itself a problem formulation

Suppose an AI helps a company perform its existing work more efficiently. The improvement is useful, but the company may still depend on the same scarce machinery, the same slow experiments, the same fragile handoffs, and the same restricted set of products. Better task execution can leave the production frontier largely unchanged.

A more consequential question is what the intelligence could build so that the next valuable result becomes easier to discover, verify, finance, manufacture, and deploy. A new algorithm can free computational capacity. A new experimental method can make a material property observable. A new process can make previously uneconomic equipment buildable. A useful capability can generate surplus that finances another bottleneck-changing capability. The desired feedback is not a larger stream of answers. It is a larger, demonstrably useful set of possible actions.

> **Transformational intelligence recognizes when the original problem formulation is unnecessarily restrictive—and discovers a fundamentally different game to play.**
>
> At civilization scale, the formulation to challenge is not only a task. It is the process by which intelligence becomes productive capacity.

This framing changes the economic question from *How much of today's work can AI perform?* to *Which realizable changes would allow people to produce outcomes that today's organization of work cannot support?* Automating an existing task can be an excellent first step. It is not a necessary ceiling on the enterprise's ambition.

Our preceding papers treated formulations, information protocols, and persistent capabilities as objects of discovery (Boucher, 2026a,b). This manuscript moves to the next level of realization: a discovered method changes a production process; an accepted product creates a possible commercial return; an authorized investment turns some return into new capacity; and that capacity supports further discovery or useful production. Each arrow can fail. A research program must identify the failure conditions rather than treat the arrows as synonyms.

1.1 An ambitious endpoint without a fictitious shortcut

The long-horizon motivation is an industrial civilization capable of using a substantial share of a star's energy flow. This is a physical scale, not a company valuation, a model score, or a claim of civilizational wisdom. The route would require energy conversion, resource extraction, manufacturing, transport, heat rejection, maintenance, and legitimate collective decisions. Intelligence could improve those processes; it cannot substitute a balance-sheet entry for their physical realization.

The relevant conceptual precedents are substantial. Transformational creativity concerns changes to the space of possible constructions (Boden, 1998; Ritchie, 2006). Endogenous-growth theory treats knowledge creation as an economic activity, while production models expose complementary inputs and the requirements of expansion (Romer, 1990; von Neumann, 1945). AI-and-growth analyses already distinguish effects on production from effects on research and emphasize that bottlenecks matter (Aghion et al., 2017; Acemoglu, 2025). Our contribution is not a new claim that innovation causes growth. It is an executable, evidence-linked formulation of *which changes unlock a productive path, how the path is financed, and where physics stops the extrapolation*.

1.2 Contributions and claim boundary

The paper makes four contributions: an operational separation of discovery, realization, surplus, and capacity; a finite bootstrap result with complete conservation records and an investment-timing counterexample; a conditional test for scalable production that does not confuse expansion with reachability; and an Alpha/SUCCESSOR research architecture tied to practical experiments and a dimensional stellar-scale ledger.

The numerical production study uses standard exact planning over investigator-supplied processes. No frontier model invents a process or performs a real transaction in the experiment. The physical calculations apply established relations to explicit hypothetical assumptions. Neither study forecasts an industrial timeline, a company's revenue, or humanity's probability of reaching a stellar-scale endpoint. Their purpose is to make the proposed causal mechanism inspectable and falsifiable.

2 Four conversions that cannot be collapsed into one score

2.1 Discovery is not yet a productive asset

A convincing proposal is not a verified capability. A capability is an executable method, together with its admissible inputs, conditions of validity, outcome evidence, costs, and recovery requirements. The artifact might be code, a tested process specification, a mathematical construction, or an experiment and its validated interpretation. Its type determines what checking is appropriate.

LLM-guided program discovery already provides real technical precedents. AlphaEvolve combines generated code, automated evaluation, and evolutionary search across mathematical and computational applications (Novikov et al., 2025). Such work motivates the first conversion, but does not make a general model's suggestions universally correct or commercially useful. Likewise, a laboratory that couples computational proposals to robotic experiments exemplifies the need to move beyond descriptions; physical characterization remains a distinct evidential obligation (Szymanski et al., 2023).

2.2 A productive asset is not automatically a profitable business

A method can be correct and socially useful while earning its owner little. Buyers may not adopt it; delivery may be expensive; alternatives may be cheaper; competition may reduce prices; or its license may not support the intended use. A revenue claim therefore needs a beneficiary, a deliverable, an adoption mechanism, and actual economic evidence. A forecast should retain its status as a forecast.

Conversely, a transfer of money can occur without creating new productive capacity. An asset-price increase, a token repricing, and an accepted useful service are not interchangeable observations. The company-level ledger should report what was delivered, what was paid, what costs were incurred, and what liabilities remain. The scientific ledger should report what improved and why.

2.3 Cash is not physical capacity

An enterprise can have enough money to order equipment but still face a long queue for fabrication, grid connection, launch, or installation. Cash purchases claims on resources through actual institutions and suppliers. It does not add kilograms, watts, machine-hours, or tested designs directly to the state of the world.

The connection between AI and electricity is already a practical systems question, not merely a distant-space concern. The IEA treats both the electricity requirements of AI and the possible effects of AI on energy operations, innovation, and security (IEA, 2025). In this paper, monetary and physical accounts remain separate even when they constrain the same decision.

2.4 Greater capacity is not identical to greater human benefit

Capacity creates options. Whether those options benefit people depends on what is produced, who can access it, which obligations are preserved, and which costs are imposed elsewhere. The purpose contract must therefore contain outcomes and material constraints beyond the system's internal growth metric. More energy use is not itself a measure of welfare; greater private wealth is not itself a complete measure of social progress.

Table 1: A conversion chain with distinct evidence at every boundary. No single receipt or benchmark score establishes the whole chain.

Conversion	Required realization	What would falsify the claim?
Idea → capability	Runnable artifact and an appropriate outcome check	The result fails on a relevant case or depends on an unreported assumption
Capability → surplus	Accepted delivery, adoption, realized receipts, full costs	No buyer, negative delivery margin, or uncounted review and maintenance
Surplus → capacity	Paid inputs, commissioned assets, resource custody	Cash is spent but the required capacity is not delivered or usable
Capacity → progress	Additional authorized outcomes without sacrificing prior obligations	Growth in a proxy while required services or affected people become worse off

The transformational opportunity can arise at any boundary. An easier-to-check artifact can improve realization. A reusable service can change delivery economics. A new manufacturing method can change what investment can buy. A different institutional arrangement may reduce unnecessary coordination costs, but changes to rights and authority require the relevant consent; they are not unilateral optimizations available to an agent.

3 Make the production network an object of discovery

3.1 State, purpose, and admissible transformations

Let the system state be

$$X_t = (\mathcal{K}_t, \mathbf{k}_t, \mathbf{m}_t, c_t, \mathbf{q}_t), \tag{1}$$

where $\mathcal{K}_t$ is validated knowledge and retained artifacts, $\mathbf{k}_t$ installed productive capacities, $\mathbf{m}_t$ inventories, c_t liquid funds, and $\mathbf{q}_t$ work in progress and pending qualifications. These coordinates have different units. They cannot be added into an undifferentiated quantity called intelligence or wealth.

A purpose contract $\mathcal{I}$ specifies required outcomes, prohibited side effects, authority, budgets, and the horizon of evaluation. Exogenous conditions z_t include demand, supplied materials, weather or solar flux, failures, and other relevant external events. A production formulation

$$F = (G, \mathcal{P}, \mathcal{M}, \mathcal{D}, \mathcal{Q}, \Pi) \tag{2}$$

contains a dependency graph G, admitted process recipes $\mathcal{P}$, information representations $\mathcal{M}$, delivery and investment rules $\mathcal{D}$, evidence-acquisition rules $\mathcal{Q}$, and an executable policy language Π. A policy chooses activities; a transformation changes a diagnosed restriction in this setup.

An ordinary sufficiently general planner can represent both choices. The distinction is not an assertion of new computability or a privileged brand of agent. It is a distinction between optimizing activities inside a fixed process and discovering or realizing a different process. Strong unrestricted planning is therefore a necessary comparator, not an opponent to exclude.

3.2 A ledger that does not let knowledge create matter

For a simple discrete production network, let $\mathbf{u}_t \geq 0$ be activity levels. A feasible schedule might satisfy

$$A_F\mathbf{u}_t + \mathbf{d}_t \preceq S_F\mathbf{k}_t + \mathbf{r}_t, \tag{3}$$

$$\mathbf{m}_{t+1} = \mathbf{m}_t + \mathbf{i}_t - C_F\mathbf{u}_t + R_F\mathbf{u}_t, \tag{4}$$

$$\mathbf{k}_{t+1} = (I - \Delta_F)\mathbf{k}_t + B_F\mathbf{u}_{t-\ell}, \tag{5}$$

$$c_{t+1} = c_t + \text{settled receipts}_t - \text{cash outflows}_t. \tag{6}$$

Here $\mathbf{d}_t$ reserves necessary service and maintenance flows, $\mathbf{r}_t$ supplies external flow resources, $\mathbf{i}_t$ gives inventory imports, R_F records physically justified recovery, Δ_F depreciation, and ℓ commissioning delay. Each row uses its own units. Monetary affordability is an additional constraint on outflows *before* later receipts become spendable.

These equations are an accounting template, not a complete industrial model. Some operations require integer decisions, nonlinear yields, uncertain outcomes, or several lead times. A real implementation must use the appropriate laws rather than force them into linear notation. Recovery cannot exceed what the operation physically makes recoverable, and no entry in R_F grants perpetual material creation.

A discovery can propose a new recipe, a better representation, a changed dependency, or an alternative investment sequence. The proposed coefficients are hypotheses until supported. A language model cannot lower A_F by editing a file and thereby claim that the factory uses less energy. The changed process must be realized and tested against the original output requirements.

3.3 Production cuts locate the right kind of problem

Proposition 3.1 (A resource cut cannot be paid away)**.** *Suppose every admitted method for producing one required output consumes at least $a > 0$ units of a non-substitutable resource, and the total available amount over the comparison is at most B. Any feasible policy produces at most B/a such outputs. Increasing computation or cash alone cannot exceed the bound while those assumptions remain true.*

Proof. Producing n outputs requires at least na units of the resource. Feasibility implies $na \leq B$. Neither extra reasoning nor a larger monetary coordinate changes this inequality unless it changes the admitted resource supply or the valid production requirement. □

The productive move is to determine whether the bound comes from a genuine necessity or an inherited method. A better process might reduce a; a commissioned supplier might increase B; a different authorized output might avoid the resource, but that would require a new purpose comparison. Pretending the resource is unnecessary is not another way around the inequality.

In coupled systems, several cuts can bind simultaneously. Improving only one may have zero effect on realized output. That is where a short-horizon policy that values only immediate output can reject the very investment needed to enter a new productive regime. The finite study below makes that interaction exact.

4 An executed study: unlock the means of expansion

4.1 A deliberately small, fully specified economy

The study is a deterministic finite planning problem, not a calibrated model of an actual factory. A seed contains $k_0 = 1$ expansion module. Each installed module supplies one unit of residual energy service and one unit of residual fabrication service per period, after the specified legacy-service reservation. These services are flows, not accumulable inventories. New modules become available in the next period. Module

construction is indivisible and must finish within one period; the main model has no multi-period work in progress or external purchases of finished modules.

The system must deliver one legacy service in each of $H = 6$ periods. The service produces a *net settled receipt* of $r = 2$ cash units at the end of each period, after its stipulated operating costs. Demand is fixed at this one service, not assumed to grow with capacity. Residual expansion flows are defined after that reservation. Within the six-period model, additional module upkeep is included in the stipulated net capacity, depreciation is zero, and there are no unmodeled imports. These are assumptions, not empirical facts.

Initially there are $c_0 = 12$ cash units and $m_0 = 20$ units of raw material. Each new module costs one cash unit and one material unit in addition to its two process-flow requirements. Material stock decreases; it does not reproduce. The objective is lexicographic: maximize installed capacity at the horizon while satisfying every service and resource constraint, then maximize remaining cash among equal-capacity plans. Capacity and cash are reported separately; the objective is not a claim that a module is worth a particular amount of money.

4.2 Two admitted projects change different requirements

The original recipe consumes $a = 2$ energy-service units and $b = 2$ fabrication-service units per new module. Project E reduces a to one; project F reduces b to one. Each project costs two cash units. Any nonempty project portfolio incurs one initial qualification period before the changed recipe is used. Project effects are supplied and assumed valid in this model; the program chooses and schedules them, but does not invent or empirically establish them.

For a fixed portfolio chosen initially, let K be its fee and d its qualification delay. The post-purchase treasury is $c_0' = 12 - K$. At period t, the integer number of newly constructed modules obeys

$$0 \le u_t \le \min\{\lfloor k_t/a \rfloor, \lfloor k_t/b \rfloor, c_t, m_t\}, \qquad u_t = 0 \text{ for } t < d. \tag{7}$$

The state updates are

$$k_{t+1} = k_t + u_t, \qquad c_{t+1} = c_t - u_t + 2, \qquad m_{t+1} = m_t - u_t. \tag{8}$$

The end-of-period receipt cannot finance construction earlier in that period. Every trace explicitly records service delivery, project costs, resource use, cash, and material.

Proposition 4.1 (An invariant bootstrap obstruction)**.** *Under Equation* (7)*, if* $k_0 < \max(a, b)$ *and no other action increases* k_t*, then* $k_t = k_0$ *for every period, regardless of the cash or material available. In particular, at* $k_0 = 1$*, the unchanged recipe and either single project cannot construct the first additional module.*

Proof. At least one of $\lfloor k_0/a \rfloor$ and $\lfloor k_0/b \rfloor$ is zero. Hence $u_0 = 0$ and $k_1 = k_0$. Repeating the argument proves the invariant. Qualification periods do not increase capacity and do not change the conclusion for a recipe whose remaining requirement is two. □

This is an exact bound on the named production formulation. It is not a universal limit on manufacturing. Allowing useful partial construction, a larger imported seed, or a different process can invalidate its assumptions. An unconstrained real-world investigation should consider those alternatives too.

4.3 The joint change produces a different reachable outcome

Neither single project increases capacity. Together they make the first additional module feasible, and each commissioned module increases the flows available for later construction. The interaction in terminal capacity is $19 - 1 - 1 + 1 = 18$ modules. A rule that accepts only individually beneficial projects can miss the pair. This is a constructed complementarity, not an estimate of how often such pairs exist in natural industries.

Table 2: Exact six-period factorial. The same seed, material stock, cash endowment, receipts, and service obligations apply. Project fees and qualification time are charged. No numerical entry is an observed industrial or monetary result.

Initial portfolio	a	b	Project fee	Final modules	Final cash
No change	2	2	0	1	24
Energy project only	1	2	2	1	22
Fabrication project only	2	1	2	1	22
Both projects	1	1	4	**19**	2

An optimal joint-project trajectory is shown below. The first period qualifies the projects while the legacy service continues. The later expansion sequence is $1, 2, 4, 8, 3$, not an indefinitely repeated doubling law.

Table 3: One optimal executed trace for the joint portfolio. Initial cash after project purchase is eight. All rows deliver the required legacy service and receive two cash units at period end.

Period	Start modules	Start cash	Built	End modules	End cash	Material left
0	1	8	0	1	10	20
1	1	10	1	2	11	19
2	2	11	2	4	11	17
3	4	11	4	8	9	13
4	8	9	8	16	3	5
5	16	3	3	19	2	2

The final cash account satisfies $12 + 6 \cdot 2 - 4 - 18 = 2$. The material account satisfies $20 - 18 = 2$. The original service count is six in every arm. New physical-like capacity is purchased with money *and* built with the available process flows and finite material; money is not substituted for those inputs.

Proposition 4.2 (The nineteen-module optimum)**.** *For the joint portfolio in the specified main model, no feasible policy can finish with more than nineteen modules. The trace in Table 3 attains this bound.*

Proof. The project leaves eight spendable cash units. Only the first five period-end receipts are available for any construction before the horizon, so at most $8 + 5 \cdot 2 = 18$ modules can be purchased. Adding the seed gives an upper bound of nineteen. The trace purchases eighteen modules and satisfies every process, material, timing, and service constraint, establishing attainability. □

The decisive change is not a forecast of faster growth. It is a verified change from "the first expansion step is unreachable" to "a feasible sequence of capacity-building steps exists."

The result also resists a misleading wealth story. The unchanged system finishes with *more liquid cash*, twenty-four rather than two. The transformed system finishes with more installed capacity. Whether that capacity later produces profitable or socially valuable outcomes requires another experiment. The nineteen-module number is not a nineteen-fold return on investment, an energy-efficiency ratio, or a company valuation.

4.4 An unrestricted comparator and a second change of game

The factorial fixes project selection at the initial procurement window to isolate the process interaction. A strong comparator should not inherit that timing restriction without examination. We therefore implement a second exact planner that may buy either project in any affordable period. Each purchase becomes effective after its qualification delay; existing production can continue while qualification is pending.

At the main twelve-unit endowment, the flexible planner also reaches nineteen modules with two cash units left. Thus the headline does not depend on excluding a better ordinary planner. Both methods are supplied the same project effects and outcome objective. The separate forward enumerator confirms the flexible planner's results.

At an initial treasury of three, however, the difference is consequential. An upfront-only system cannot purchase both two-unit projects and ends with one module. The flexible planner first delivers the legacy service, increasing cash from three to five. It then purchases both projects for four, qualifies them, and constructs $1, 2, 4, 2$ modules over the remaining four periods. It finishes with ten modules and two cash units.

Table 4: Investment timing is itself a formulation. The flexible comparator uses ordinary exact planning; this is not evidence of learned financial reasoning or a real financing transaction.

Initial treasury	Best upfront-only final modules	Flexible-procurement final modules
3	1	10
12	19	19

The second result matters conceptually: a declaration that an investment is impossible can be an artifact of when the model insists it must be funded. Conversely, future receipts must actually be available and the delay must leave enough useful time. The flexible planner does not spend forecast receipts early; it waits for the modeled settlement event. No statement of infeasibility should be transferred from the upfront-only formulation to every permitted financing sequence.

4.5 Controls that stop the attractive extrapolation

The companion evaluates ten boundary cases, a fifty-four-configuration factorial sweep, and three flexible-procurement initial treasuries. The sweep covers horizons $2, 4, 6$, seeds $1, 2$, cash endowments $6, 12, 18$, and material stocks $0, 6, 20$. All 216 portfolio results agree with separate forward reachability; their traces satisfy 12,528 recorded semantic checks. Parameter combinations are coverage tests, not independent real-world discovery campaigns.

The source model is complete, noiseless, and deliberately favorable to interpreting the mechanism. There is no physical calibration of the module, price, project effect, or period. The program searches supplied alternatives; it does not discover a manufacturing breakthrough or establish a customer market. Exactness removes numerical ambiguity *inside* the model, not uncertainty about the model's relevance.

5 When can a productive loop actually compound?

5.1 Productive closure is not thermodynamic closure

A production network is productively closed over a declared activity set when every required *produced* input has an admitted source, with adequate capacities, inventories, qualifications, and maintenance. It may remain open to external energy, raw material, information, people, and authorized trade. Calling such a network self-expanding must never conceal those inputs.

Space-industry bootstrapping has a relevant prior literature: sub-replicating systems can be analyzed as

Table 5: Selected exact controls. "Best" uses the stated capacity-first objective and then remaining cash. Upfront-only timing applies except where stated.

Changed assumption	Final modules	Supported interpretation
Eighty cash units, no recipe change	1	Cash alone does not cross the process cut
No material stock, best portfolio	1	Choose no projects; matter is not created
No recurring receipts, best portfolio	9	Construction is limited by the finite treasury
Only six material units, best portfolio	7	The raw-material ceiling becomes binding
Qualification takes all six periods	1	Choose no projects; no time to realize the gain
Fabrication project not permitted	1	Authority is not relaxed to obtain a higher score
Both projects leave requirements unchanged	1	Decline ineffective projects
Seed already contains two modules, no change	19	The original obstruction was seed-dependent

candidates for progressively producing more of their own inputs (Metzger et al., 2013). We use this as a modeling precedent, not as evidence that any proposed launch mass, replication rate, or completion date has been demonstrated. A bill of materials that omits a critical imported component is not closed because its omitted mass fraction is small.

5.2 A finite-scope expansion certificate

Let $\mathbf{v} \succ 0$ be a reference vector of productive capacities, including every essential capacity rather than only the fastest-growing sector. Consider a qualified production schedule that can be scaled over a stated finite range. Extra resources may be left idle; an excess must not force a violation.

Theorem 5.1 (Conditional constructive expansion)**.** *Suppose that for every scale λ in a declared domain $[\lambda_0, \lambda_{\max}]$, a permitted schedule starting from at least $\lambda\mathbf{v}$ (i) preserves the required outcomes, (ii) pays maintenance, construction, validation, and all other counted costs, (iii) has the needed material and external-flow supply, and (iv) ends after a fixed cycle duration with at least $\gamma\lambda\mathbf{v}$, where $\gamma > 1$. If a feasible bootstrap reaches $\lambda_0\mathbf{v}$, then the schedule constructs at least $\lambda_0\gamma^n\mathbf{v}$ after n cycles, for every n whose preceding cycle scales remain in the domain and whose supply assumptions continue to hold.*

Proof. The bootstrap supplies the initial lower bound. Apply the admitted schedule at λ_0. It preserves obligations and returns at least $\gamma\lambda_0\mathbf{v}$. Repeating the same argument at scale $\gamma\lambda_0$, leaving any excess idle if necessary, proves the claim by induction. The induction stops when a required scale, supply, or qualification leaves the stipulated domain. □

The theorem is deliberately conditional. It gives a constructive sufficiency test, not an algorithm for discovering the schedule or a claim that all industries satisfy it. A favorable eigenvalue in an abstract reproduction matrix is insufficient when the matrix omits shared energy, a fabrication queue, commissioning delays, material imports, or obligations. Likewise, a long-run expansion schedule need not be reachable from a small seed. The bootstrap and the scalable cycle require different certificates.

A finite stock of material immediately limits the duration of the claim. In the main experiment, $m_0 = 20$ and one unit is embodied in each additional module; hence no policy can exceed twenty-one installed modules without new material. The later cash-limited value nineteen is even smaller. Neither theorem nor software turns a finite reservoir into indefinite exponential growth.

5.3 A moving bottleneck, not a permanent growth constant

Once one restriction is relaxed, another can become decisive. The main trajectory first encounters an indivisible process threshold, then qualification time, then cash, with material close behind. At larger scales, maintenance, control, energy collection, heat rejection, and supply-chain closure can become limiting. An intelligent expansion policy should diagnose the *next* bottleneck rather than extend the previous growth rate by assumption.

The correct target is a changing reachable set under a purpose contract. Let

$$\mathcal{R}_{\mathcal{I}}(X_0, F, B, H) \tag{9}$$

contain the outcome-and-capacity vectors attainable by permitted trajectories from the stated initial condition within the declared resource boundary and horizon. A useful transformation supplies a new admissible trajectory, pays its construction cost, and improves a relevant part of this set. An observed gain against implementations is weaker than a proved exclusion from the old set. The finite bootstrap invariant supplies an exclusion; most natural workloads will require empirical comparisons instead.

6 The enterprise: turn validated usefulness into reinvestible surplus

6.1 What could make an owned intelligence system highly valuable?

An enterprise that repeatedly produces difficult, useful, independently checkable results could possess a valuable productive asset. The commercially relevant system is not only a model checkpoint. It can include validated methods, operational tooling, trusted evidence, customer relationships, reproducible delivery, and the ability to discover what is worth building next. Which components can be owned or licensed depends on their provenance and terms; an open-source foundation does not imply exclusive ownership of its shared code.

Knowledge can be reusable while equipment and human attention remain rival inputs. This distinction is central to endogenous-growth models (Romer, 1990). The enterprise opportunity is to combine reusable knowledge with dependable realization, not assume that ideas command unlimited prices. Competition, diffusion, and diminishing marginal demand can transfer benefits to customers or society while reducing the originating firm's margin.

For a delivered capability j, distinguish customer value V_j^{user}, realized receipts R_j, and base discovery-and-delivery cost C_j^{base}, excluding the obligations listed separately below. A simple firm-level operating surplus is

$$\Pi_j = R_j - C_j^{\text{base}} - C_j^{\text{warranty}} - C_j^{\text{maintenance}} - C_j^{\text{other}}, \tag{10}$$

where categories must be defined to prevent double counting and include all applicable obligations. The user's value is neither asserted equal to the price nor automatically known. A price, an invoice, a settled payment, and free cash available for investment are distinct records.

6.2 The scientific and commercial tests should reinforce one another

A useful initial offer is a bounded mission with an outcome contract, an incumbent, alternative formulations, a decisive experiment, and a deliverable. The buyer funds specified investigation rather than a guaranteed breakthrough. The result may be a better method, a justified refusal to change, or a finding that the expensive idea is not worth implementing.

Commercial confirmation then asks whether the buyer adopts the result, whether it performs over the agreed period, and whether receipts exceed the complete cost of delivering it. A valid result with negative margin is not a profitable service. A profitable service with no independently grounded improvement is not

evidence for a scientific transformational claim. The strongest case satisfies both tests and records them separately.

Complementary organizational and intangible investments can be substantial before a technology's benefits appear; the productivity J-curve literature is a useful warning against counting only the new tool's invocation cost (Brynjolfsson et al., 2021). Review, integration, monitoring, failed candidates, and founder or operator time belong in the accounting. A system that scales output by silently scaling unpaid human rescue has not demonstrated the claimed economic mechanism.

6.3 Reinvestment changes the future opportunity set

Let s_t be realized, unencumbered surplus after required reserves. An authorized investment decision allocates it among maintaining existing services, acquiring evidence, developing capabilities, buying or building capacity, distributing returns, and retaining liquidity. A conceptual choice is

$$a_t \in \mathcal{A}_{\mathcal{I}}(X_t, s_t), \qquad a_t \in \arg\max_a \mathbb{E}\big[U_{\mathcal{I}}(\tau_{t:H}) \mid X_t, a\big], \tag{11}$$

with hard constraints outside the scalar utility and noncommensurable objectives reported separately. This is a decision specification, not an oracle for unknown demand or physical outcomes.

The transformational move is to ask which investment changes what future investment can achieve. A cheaper experiment may avoid financing a useless build. A standard interface may enable independent suppliers. A new process may lower the seed needed to produce the next machine. The effect is contingent on delivery and adoption, not automatic reinvestment into the same kind of asset.

An enterprise valuation would require assumptions about future cash flows, competition, risk, financing, and distributions. None is estimated here. The manuscript therefore supplies no valuation, price target, return forecast, or claim that present token balances represent future physical capacity. Its ambition is a productive mechanism, not a speculative multiple.

6.4 Evidence-conditioned value rather than multiplied hype

A promising chain has several conditional events: the idea is valid, implementation works, a customer adopts it, cash is received, and the next capacity is commissioned. If E_i are those events, then

$$\mathbb{P}\left(\bigcap_{i=1}^{n} E_i\right) = \prod_{i=1}^{n} \mathbb{P}\left(E_i \mid \bigcap_{j<i} E_j\right). \tag{12}$$

This is the probability chain rule, not an assumption of independent success. Multiplying unrelated marginal success rates can be misleading when one common failure affects several stages. Where the probabilities are unmeasured, the correct representation is uncertainty and an experiment plan, not an invented precision.

Confidence must also survive system changes. A method validated on one workload may not transfer; a price can fall; a supplier can fail; a contract can change. Retain the capability's evidence and its invalidation conditions. The system should retire, repair, or requalify when the premise of its return disappears.

7 The stellar-scale boundary: power, matter, heat, and time

7.1 Specify the physical endpoint

A stellar-scale, or Type II, energy capability concerns access to power on the scale of a star, not the size of a monetary account (Kardashev, 1964). Here we study absorption fractions up to order unity, keeping absorbed, converted, and delivered power distinct. For numerical transparency we use the IAU nominal solar luminosity

$$L_{\odot}^{\mathrm{N}} = 3.828 \times 10^{26}\ \mathrm{W}, \tag{13}$$

an exact nominal conversion constant, not a claim that the Sun's instantaneous luminosity is invariant (Prša et al., 2016). Let $f \in (0, 1]$ be the fraction of stellar radiation absorbed by the infrastructure, and let η be an assumed conversion efficiency. Then

$$P_{\text{abs}} = f L_{\odot}^{\text{N}}, \qquad P_{\text{electric}} = \eta f L_{\odot}^{\text{N}}. \tag{14}$$

Electrical power is not automatically useful delivered service: distribution, storage, computation, conversion, and maintenance consume part of it. A civilization can increase useful outcomes without maximizing either expression.

A distributed collection and industrial architecture is a hypothesis to investigate, not a completed design in this manuscript. The thermodynamics of radiation-powered stellar infrastructure already has a substantial technical treatment (Wright, 2023). Our purpose is to expose what an AI-led expansion claim must account for, not replace detailed thermal, orbital, or structural engineering with a geometric sketch.

7.2 Collection area is a material commitment

For ideal nonoverlapping, fully absorbing collectors all at radius r, the equivalent projected collection area is

$$A_{\text{col}} = 4\pi r^2 f. \tag{15}$$

This is an ideal geometric relation. Orientation, gaps, shadowing, optical losses, finite lifetimes, and orbital feasibility require separate treatment. It is not a stable-swarm construction certificate, especially near complete coverage. With an assumed collector areal mass μ_{col},

$$M_{\text{col}} = \mu_{\text{col}} A_{\text{col}}. \tag{16}$$

The calculation uses 1 au = 149 597 870 700 m exactly under the IAU definition (IAU, 2012). Areal mass is a hypothetical engineering input, not an established universal material limit.

7.3 Useful computation does not erase waste heat

Consider a steady boundary in which all absorbed stellar power, including energy temporarily converted into useful work, ultimately dissipates locally. Then $P_{\text{heat}} = P_{\text{abs}}$. It would be incorrect to count only $(1 - \eta)P_{\text{abs}}$ as the eventual heat load when the electrical output is also used and dissipated inside that boundary. Exported work or energy can alter this balance, but must be explicitly recorded rather than silently discarded (Wright, 2023).

An ideal grey radiator with emitting area A_{rad}, emissivity ϵ, temperature T, and radiative background T_b satisfies

$$A_{\text{rad}} \geq \frac{P_{\text{heat}}}{\epsilon\sigma(T^4 - T_b^4)}. \tag{17}$$

We use $\sigma = 5.670374419\ldots \times 10^{-8}$ W m^{-2} K^{-4} from CODATA (NIST, 2022). The expression is an ideal lower bound neglecting parasitic loading, mutual irradiation, limited view factors, transport losses, and other engineering complications. Emitting area is not the same quantity as projected collector area. The calculation does not prove that both can be arranged together as assumed.

At one-percent absorption, gross electrical power under the assumed efficiency is 1.1484×10^{24} W. At a hypothetical collector areal mass of 0.1 kg/m^2, the collector sheets alone weigh about 2.812×10^{20} kg. With a separate hypothetical radiator areal mass of 1 kg/m^2, the 300 K emitting-area lower bound corresponds to about 9.260×10^{21} kg of radiator material. These are sensitivity calculations, not component specifications or a total-system bill of materials. Shared-function surfaces, supporting structure, power conversion, shielding, and transport would change the accounting.

Table 6: Computed dimensional ledger, not an engineered system or forecast. Assume $r = 1$ au, $\eta = 0.30$, $\epsilon = 0.90$, $T_b = 3$ K, and local dissipation of all absorbed power. Rounding is to four significant figures.

Absorbed fraction f	P_{abs} (W)	A_{col} (m^2)	A_{rad} (m^2), 300 K	A_{rad} (m^2), 600 K
10^{-12}	3.828×10^{14}	2.812×10^{11}	9.260×10^{11}	5.788×10^{10}
10^{-6}	3.828×10^{20}	2.812×10^{17}	9.260×10^{17}	5.788×10^{16}
10^{-2}	3.828×10^{24}	2.812×10^{21}	9.260×10^{21}	5.788×10^{20}
1	3.828×10^{26}	2.812×10^{23}	9.260×10^{23}	5.788×10^{22}

Raising the assumed radiator temperature from 300 K to 600 K reduces its ideal area by nearly sixteen because of the fourth-power law. That is not a free solution: materials, conversion efficiency, operating hardware, and lifetime may change. A transformational research program should explore these coupled design choices while keeping the complete outcome and resource requirements visible.

7.4 A physical construction-rate bound

Let e_m be embodied energy per kilogram of qualified installed infrastructure, and let P_{build} be the power actually available for construction after other demands. Let the following rates denote usable material feed, fabrication, installation, and qualified acceptance, all in kg/s. A necessary rate constraint is

$$\dot{M}_{\text{installed}} \leq \min\left\{\frac{P_{\text{build}}}{e_m}, \dot{M}_{\text{feed}}, \dot{M}_{\text{fab}}, \dot{M}_{\text{install}}, \dot{M}_{\text{qualified}}\right\}. \tag{18}$$

Additional limits can lower the true rate further; the displayed minimum is a necessary upper bound, not an achievable-throughput guarantee. A program that claims enormous growth should state how each term expands and how the expansion itself is funded, powered, fabricated, and checked. More planning tokens do not directly increase any of the physical rates.

Recycling reduces new input requirements but does not create net material. For an initial feed mass M_0 repeatedly processed with recovery fraction $0 \leq \rho < 1$ and no additional imports, the simple geometric model bounds cumulative gross processed mass by $M_0/(1-\rho)$. The retained or installed net mass still cannot exceed the available material mass without new input. Gross throughput, recovered inventory, and new construction must not be counted as three independent resource sources.

7.5 Doublings are not dates

For an arbitrary one-terawatt initial absorption capacity, the idealized number of doublings to the nominal solar luminosity is

$$n = \log_2\left(\frac{3.828 \times 10^{26}}{10^{12}}\right) \simeq 48.44. \tag{19}$$

Reaching one percent corresponds to about 41.80 such doublings. These are scale ratios. Assigning a constant time to each doubling would assume away the research, manufacturing, maintenance, supply, and coordination problems this paper makes central. No calendar forecast follows from this arithmetic.

The meaningful path is a succession of verified capacity regimes. A first digital result may lower the cost of computation or engineering. A later physical result may increase energy or fabrication throughput. Subsequent systems may rely on qualified off-world material and industry rather than shipping every component from Earth. Each transition requires its own evidence; the present finite simulation is not a calibration for any of them. The bridge to stellar-scale capability is the sequence of realized transitions, not an extrapolated software growth curve.

8 SUCCESSOR Ω: an implementation of the research loop

8.1 The project is the starting platform, not a distant capability claim

The named implementation target is the open-source **AGI Alpha Agent** project:

https://github.com/MontrealAI/AGI-Alpha-Agent-v0

Find your Alpha. Build your SUCCESSOR Ω.

Here *Find your Alpha* means identify a valuable, removable restriction and the evidence needed to challenge it. *Build your SUCCESSOR* Ω means construct and retain a useful capability that improves what the next mission can actually achieve. The successor is not merely a renamed model, a copied agent, or a longer memory. It is an owned or appropriately licensed, portable system of capabilities, evidence, and conditional operating rights.

The repository snapshot at commit `9c7484d...`, inspected for this manuscript, presents Computer Use, Level 3 strategy development, and Level 4 problem reframing in one flagship journey. Its guide distinguishes free browser preparation from separately configured native execution, and distinguishes proposals from executed patches, verified mathematical results, and observed economic improvements (MONTREAL.AI, 2026a,b). These distinctions should be preserved. The current guide is not evidence of an autonomous physical industry or a qualified realization of every component proposed here.

8.2 One outcome, several possible games

A user should begin with an authorized outcome and approved public inputs, not a requirement to supply the best problem formulation in advance. Alpha should investigate whether the next action ought to be a better solution, a different representation, an experiment, a changed workflow, a reusable artifact, a commissioned capacity increase, or a decision not to proceed.

Computer Use provides one means of interacting with tools. It should expand access, not force every solution to reproduce an unnecessary sequence of clicks. Where a permitted API, solver, compiler, or different workflow better serves the outcome, discovering that change is part of the research target. New execution capabilities require their own adapters and qualification; they are not implied by a browser's ability to display a proposal.

The five functions discussed in the preceding project work fit one loop: opportunity discovery selects a consequential restriction; transformational discovery constructs an alternative; experiment invention makes its mechanism testable; capability construction retains what works; and meta-strategic evaluation tests changes to the process that discovers all four. They need not be five separate agents or products.

8.3 A capability-to-capacity record

Every proposed improvement should retain a typed record that binds the discovery to the claimed conversion. In particular, a paper, a simulation, a customer acceptance, a payment, and a commissioned machine must not share one undifferentiated success label.

The finite study supplies a prototype of only the simulation-and-checking part: a bounded model, project portfolio, schedule, conservation trace, and separately checked terminal result. It is not a payment instruction or equipment order. The physical ledger supplies dimensional constraints, not an approved construction plan.

8.4 A concrete next software increment

A proposed versioned `production-network/v1` adapter would admit bounded models, alternative recipes, scenario sets, objective contracts, and independently checkable schedules. It should extend the existing

Table 7: Proposed record for an Alpha/SUCCESSOR mission. The fields are design requirements, not a claim that the repository already implements this schema.

Field group	Required meaning
Purpose and authority	Beneficiary, desired outcomes, nonnegotiable requirements, input rights, spending and execution permissions
Restriction and alternative	Named incumbent, diagnosed bottleneck, changed representation/process/investment rule, and what must remain invariant
Evidence	Source and candidate versions, actual observations, independent check, fresh comparison, failure cases, and uncertainty
Value realization	Accepted deliverable, adoption status, actual receipts, complete costs, reserves, and unsatisfied liabilities
Capacity conversion	Resources ordered versus received, installed capacity, qualification state, lead time, maintenance, and residual constraints
Retention and succession	Supported contexts, portable artifact, provenance, invalidation triggers, review, rollback, and new mission tests

mission history rather than create a disconnected application. A configured model may propose a candidate formulation as data; only admitted execution mechanisms run it. The independent checker evaluates the original outcomes and constraints, not the candidate's preferred score.

The adapter should return at least four separate artifacts: a candidate method, actual test outputs, a complete resource ledger, and a claim record stating whether the result is a proposal, a modeled improvement, an observed operational gain, or a qualified retained capability. Model usage and human review time must remain explicit. A fixed-name model or another agreeing judge does not establish competence or independence.

The first natural mission should remain much nearer than space industry. A useful candidate is a public computational service with expensive repeated work and independently checkable outputs. The agent may discover a different algorithm, representation, or execution schedule. The service is then evaluated under an actual resource limit and a relevant workload, with all required outputs and accuracy conditions preserved. A scheduling mission should use a strong established solver, such as an appropriate OR-Tools configuration, and allow the unrestricted agent access to it as well (Google, 2026).

After a verified improvement, the next test should *use* the released resource: execute additional accepted work or commission a measured capacity increment within the stated total budget. Forecast savings do not count as an installed successor. The experiment should also include a future case where reuse is inappropriate, so that retention is tested as conditional judgment rather than universal replay.

8.5 Alpha, Jobs, and Fleet without conflating their evidence

The proposed division of responsibility is straightforward. Alpha organizes discovery and construction; Jobs commissions bounded work and review under admitted terms; Fleet supports qualified execution and evidence transport. The primary user experience should remain one coherent mission, while each component retains a precise authority boundary.

AGIJobManager's USDC route and the Alpha/Jobs v0 AGIALPHA route are distinct in the inspected documentation (MONTREAL.AI, 2026a,b). Neither should be treated as a financial guarantee of a result. A research request with no spending allowance must not become a paid model call, token transfer, or physical procurement simply because a planner judges the opportunity attractive. Treasury actions, deployments, and physical operations require separate authorization and qualification.

This manuscript creates local research artifacts only. It makes no repository write, release, deployment,

wallet transaction, or external operational change. Its implementation proposals preserve the original papers, histories, existing functionality, and separately owned integration and remediation work.

8.6 The discovery process can become a better successor

Level 3 meta-strategic intelligence concerns how alternatives are generated, investigated, compared, and combined. Level 4 transformational intelligence concerns the formulation that determines what can be represented, measured, built, or achieved. A successor can improve either, but the evidence differs.

Let η be a discovery-and-investment policy. A campaign begins with a fixed initial capability library, task environment, resource endowment, and permission set Z. After a development budget it returns a system whose performance is measured on untouched terminal conditions. Its value can be specified as

$$\mathcal{V}(\eta) = \mathbb{E}_Z\left[U_I(\tau_{\eta,Z}^{\text{terminal}}) - C_{\eta,Z}^{\text{discovery}}\right], \tag{20}$$

only where costs and utility have a justified common scale; otherwise they form separate constrained objectives. Operating costs inside terminal utility must not be subtracted twice.

An improvement to η needs multiple independent complete campaigns, not just one successful descendant. Cross old and new policies with old and new starting libraries when needed to separate better discovery from a better inheritance. DGM and Hyperagents provide direct precedents for editable agents and editable improvement procedures (Zhang et al., 2025, 2026); they are comparison targets, not reasons to attribute unmeasured civilizational effects to a self-modification loop.

9 An empirical program that can fail informatively

9.1 The strongest next result is an observed conversion

The next substantive result should begin with an independently selected natural problem whose useful reformulation is not supplied as the answer. It should show a diagnosed restriction, a realized alternative, a strong comparison, an accepted outcome, and a measured conversion of released resources or receipts into another useful capability. Demonstrating every step on one bounded mission is stronger than presenting separate favorable anecdotes for each step.

This does not require the first mission to be scientifically unprecedented. Rediscovering a known method from appropriate starting information can demonstrate useful discovery. Scientific novelty, outcome validity, commercial adoption, and repeatable economic advantage are different claims. The report should identify which ones survive testing.

9.2 Baseline access is part of the experiment

Compare the incumbent, a strong fixed-formulation method, a strong unrestricted agent explicitly allowed to reformulate and reuse, and a relevant domain-engineered system. Where appropriate include automated agent or workflow search such as ADAS (Hu et al., 2024). Match model access, permissions, starting information, and total resources; count the external engineering used to supply candidate recipes or instruments.

A factorial test should isolate the changed formulation from its compatible implementation. Giving a baseline the discovered representation can reveal whether discovery or execution is responsible for the gap. Removing the structural mechanism while retaining comparable support tests attribution. The flexible-procurement control in Section 4.4 illustrates why a seemingly innocuous initial restriction can alter the conclusion.

For physically mediated work, use independently measured final states and production records. For digital computer work, environments with inspectable states provide a relevant precedent in OSWorld (Xie et al.,

2024). Neither a fluent transcript nor a successful script on a synthetic fixture is evidence that a general model delivered a real-world outcome.

9.3 Evidence quality and evaluator revision

Candidate selection and confirmation should be separated. Freeze the finalist, model/tool configuration, outcome contract, comparison rules, and cost boundary before fresh confirmation. Adaptive reuse of evaluation data can invalidate ordinary inference (Dwork et al., 2014); valid sequential methods must match their sampling and stopping assumptions (Howard et al., 2021). Service episodes, customers, repositories, and complete discovery campaigns can be different statistical units.

An evaluator can itself preserve an unnecessary restriction. A metric that counts old task completions may penalize a construction that eliminates those tasks while preserving their purpose. The replacement instrument must be validated separately against the intended outcome; both incumbent and candidate should then be compared on the common new basis. A candidate does not earn success by selecting its own easier judge.

For a chain of evidence claims with valid error bounds α_i, a union bound gives a total error probability no greater than $\sum_i \alpha_i$, without requiring independence. But a mathematical bound on sampling error does not include unknown model inadequacy, missing obligations, or unmeasured economic risk. An unknown error must not be inserted as zero. Hashes bind records to bytes; signatures identify configured principals; neither establishes semantic truth or organizational independence.

9.4 Six falsifiable hypotheses

H1: useful production reformulation. On unfamiliar natural workloads, the system discovers and realizes changes that improve purpose-grounded outcomes or full lifecycle cost against strong unrestricted baselines. The claim fails when improvement relies on extra permissions, uncharged engineering, or changed obligations.

H2: bootstrap-aware selection. The system identifies complementary changes or alternative investment sequences that make a previously unreachable productive step feasible. A comparative-advantage claim fails when an equally informed ordinary comparator reaches the same outcome more cheaply. Independently recovering a known feasible path supports a narrower discovery claim; a proposed seed that cannot support the first step supports neither.

H3: evidence creates better investment decisions. A selected experiment changes a material build, stop, or reuse decision and improves its outcome after experiment cost is included. Generating more hypotheses or documents is not sufficient. Uninformative observations and unaffordable experiments must sometimes be rejected.

H4: surplus becomes measured capacity. Accepted work produces realized net surplus or verified released resources, and an authorized reinvestment produces additional usable capacity without impairing earlier obligations. The claim fails if it records only forecast revenue, asset repricing, unused credits, or uncommissioned equipment.

H5: persistence survives change. Retained capabilities remain valid within their stated future scope and detect cases requiring recovery, requalification, or retirement. The claim fails when reduced work merely postpones errors, shifts costs to others, or relies on an uncounted external dependency.

H6: discovery improves across campaigns. A changed discovery policy produces more valuable verified transformations on new mechanism families under matched initial resources. The claim fails when the gain comes only from inheriting an already superior library, repeated evaluation exposure, or increasing human rescue effort.

9.5 A staged evidence ladder rather than a launch date

Table 8: A possible sequence of research milestones. These are proposed evidence requirements, not claims of completed work or dated forecasts.

Stage	What must be built	Evidence needed to advance
Digital capability	A useful algorithm, tool, or workflow on approved public inputs	Executed artifact, independent checker, fresh comparison, complete cost record
Commercial realization	A bounded repeatable service	Adoption, accepted delivery, realized receipts, positive lifecycle economics where claimed
Capacity formation	A commissioned computational or engineering increment	Actual usable throughput, reliability, maintenance, and no sacrificed obligations
Physical production	A qualified process or component improvement with partners	Independent measurement, repeatability, materials/energy ledger, safety and operating approval
Off-world industry	A bounded modular supply and construction demonstration	Launch and local production evidence, realistic input closure, repair, and heat accounting
Stellar-scale systems	Successive resource-accounted capacity regimes	No missing critical inputs; verified scale transitions rather than extrapolated doubling assumptions

The ladder is not a demand to centralize every stage inside one company. A company can retain its core assets while partnering, licensing, purchasing, or creating separately governed ventures for capabilities it does not possess. The scale of the objective increases the need for interoperable evidence and legitimate coordination; it does not imply that one owner should acquire unilateral authority over a civilization.

10 Institutions and scope are part of productive capability

An advanced system may produce valuable proposals at a rate that existing review or coordination cannot support. The transformational response is to investigate better evidence, modular interfaces, and proportionate review, not remove oversight merely because it slows the desired outcome. A valid new process should make required checks easier or more informative while preserving the outcome they protect.

The purpose and authority contract must distinguish owner-approved digital actions from external commitments, financial execution, physical intervention, and decisions affecting people outside the project. A model can propose a change to this contract. It cannot make the change authoritative by embedding it in a generated plan. Exploration may be broad within an isolated research setting; operational authority remains scoped and separately granted.

Productive capability should also be portable and recoverable. Retain versioned artifacts, evidence, dependencies, and independently held checkpoints. Support interruption, explicit budget limits, failure reporting, rollback where meaningful, and a retirement path. Physical commitments can be irreversible; reversibility must be demonstrated rather than assumed from software rollback. A stopped controller does not undo material already processed or equipment already launched.

A for-profit enterprise and broad human benefit are not mutually exclusive objectives, but neither guarantees

the other. User value, reliable service, competitive access, labor effects, ecological effects, and private returns deserve separate measurement. Expanding options is the technical ambition; legitimate human choice determines which options are exercised. The system should not infer permission to maximize energy capture or institutional control from a request to improve a useful service.

These are enabling requirements for realizing the vision at scale, not a substitute for invention. A governance checklist cannot discover a better process. A brilliant process cannot grant itself the authority to operate. The architecture must support both invention and accountable realization without confusing their roles.

11 Relation to prior work and remaining scientific gaps

Changing conceptual and computational spaces. Boden and Ritchie provide conceptual foundations for transformational creativity (Boden, 1998; Ritchie, 2006). Our preceding manuscripts made purpose-preserving formulation change and future-valid capability explicit (Boucher, 2026a,b). DreamCoder supplies a technical precedent for learning abstractions that alter program search (Ellis et al., 2020). The present increment is the connection between a changed formulation, a reachable capacity-building path, actual surplus, and physical production constraints.

Economic growth and complementary investment. Romer models technological change as a result of economic activity, while von Neumann's production framework analyzes expansion with interdependent processes (Romer, 1990; von Neumann, 1945). AI-growth research already examines both production automation and research effects, and task-based macroeconomic work stresses the importance of the actual exposed activities and their complementarities (Aghion et al., 2017; Acemoglu, 2025). The productivity J-curve highlights the required intangible complements (Brynjolfsson et al., 2021). Our finite model is not an alternative calibrated macroeconomic theory. It isolates an exact bootstrap and financing interaction and refuses to equate capacity counts with GDP or private valuation.

Automated invention and self-improvement. AlphaEvolve, ADAS, DGM, and Hyperagents establish direct comparison families for generated algorithms, agent designs, and editable improvement procedures (Novikov et al., 2025; Hu et al., 2024; Zhang et al., 2025, 2026). Their existence makes it inappropriate to present "AI improves a strategy" as an untouched field. Our empirical question is whether a complete system can discover a restriction, build a different process, realize its value, and improve the next campaign under full resource accounting. No frontier-model superiority result is reported here.

Physical realization and stellar infrastructure. Autonomous laboratories provide examples of connecting computational proposals to experimental action (Szymanski et al., 2023). Space-industry bootstrapping studies and radiation thermodynamics provide relevant models of production closure and energy processing (Metzger et al., 2013; Wright, 2023). These sources do not establish that an AI enterprise has a demonstrated route to an order-unity stellar energy share. Our physical ledger exposes the dimensions that such a route must satisfy and the measurements that would be needed to justify advancing it.

11.1 What remains unknown

The finite economy assumes known deterministic recipes, one-period indivisible construction, stipulated net capacities, fixed end-of-period receipts, and no depreciation within its short horizon. The flexible control relaxes investment timing but not every industrial assumption. In particular, multi-period assembly, trade in finished assets, uncertain projects, demand elasticity, and shared upstream dependencies require richer models. A successful synthetic trajectory does not establish that analogous natural recipes can be found or built.

The separate checker and forward enumerator reduce dependence on the constructor's implementation. They are not external institutional reviews, and both trust the source model. The dimensional physics calculations use hypothetical efficiencies, temperatures, and areal masses. They are necessary-accounting examples, not a validated orbital architecture, qualified component design, or a deployment-safety assessment.

The broadest uncertainty is discovery: whether a system can identify and realize useful transformations in unfamiliar environments, with imperfect models, realistic costs, changing institutions, and competing alternatives. The paper specifies how to test that possibility. It does not replace such tests with a forecast of exponential capability, automatic wealth, or an assured civilizational destination.

12 Conclusion: build the system that makes the next frontier reachable

The original problem can be restrictive at several levels. A task can demand unnecessary reconstruction. A workflow can preserve avoidable dependencies. A business can mistake task volume for productive value. An investment model can require all funding before the revenue that would make the project possible. A civilization-scale vision can confuse financial means with physical realization.

Transformational intelligence asks which of these restrictions is actually binding and removable, constructs a different setup, and demonstrates that the change serves the intended purpose. The finite study supplies a precise illustration: two separately ineffective changes unlock a capacity-building sequence; a separate change in investment timing unlocks that sequence from a smaller treasury. The same study exposes the limits: material, cash, permission, qualification time, and initial seed conditions can stop the apparent compounding.

The link to AGI Alpha Agent and SUCCESSOR Ω is operational. Discover a valuable restriction. Invent the experiment that tests it. Build an executable alternative. Demonstrate its useful outcome. Retain the capability and its conditions. Convert actual surplus or released resources into a measured next capacity. Then test whether the successor becomes better at choosing and realizing the next transformation.

The most consequential output of intelligence may be a better system for making its next discovery realizable.

Do not only optimize production inside today's frontier. Discover and construct the conditions under which a different frontier becomes reachable—with the evidence, resources, and authority that make the advance real.

Research status and preservation of earlier work

This manuscript contains analytic arguments, an executed exact synthetic production study, a separately implemented checker and reachability comparator, dimensional physical calculations, and an implementation and evaluation blueprint. The companion passes 83 tests. It makes no frontier-model or network calls, generates no real revenue, constructs no physical assets, and makes no repository or on-chain changes. No external peer review or general transformational-intelligence achievement is claimed. The earlier papers and their public links remain unchanged. Primary literature and the named repository snapshot were checked on October 19, 2026.

A Reproduction and interpretation of the executable artifact

A.1 Run the complete study

The companion uses the Python standard library and was checked with Python 3.13.5. It makes no network calls, reads no credentials, and performs no remote writes. From the extracted project directory:

```
python experiments/run_study.py
python -m unittest discover -s tests -v
python experiments/check_artifact.py results/capacity_certificate.json
```

The first command regenerates all six result JSON files. The second runs 83 tests, including exact comparisons, negative controls, resource-ledger mutations, finite parameter coverage, dimensional calculations, and bounded artifact input checks. The third independently verifies the saved main trace and recomputes its optimality within the fixed portfolio. Use `-trace-only` to check feasibility without making the optimality claim.

The artifact command accepts only bounded inert JSON. It rejects duplicate keys, nonfinite constants, unknown schemas, malformed quantities, false scope labels, and accounting violations. It does not execute model-produced code. These checks are a narrow local input boundary, not a general-purpose security sandbox. A valid record remains conditional on its supplied model and the trusted checker.

The standalone `Beyond_the_Production_Frontier.tex` embeds its bibliography and all tables. It requires no external figures or BibTeX pass. Compile with `latexmk -pdf`, or run `pdflatex` twice. Numerical tables are generated by the separate Python study; LaTeX is not the experiment engine.

A.2 Files and their evidential roles

File	Role
`experiments/production.py`	Fixed-portfolio source model, exact recursive optimizer, and complete main trace
`experiments/checker.py`	Separately written raw-record checker and forward reachability comparators; does not import the optimizer
`experiments/adaptive.py`	Flexible-procurement exact planner with explicit purchase and qualification dates
`experiments/stellar.py`	Dimensioned area, power, mass, heat, and scale calculations under declared assumptions
`experiments/run_study.py`	Deterministic result generation and cross-checks
`experiments/check_artifact.py`	Bounded command-line validation of an inert fixed-portfolio trace
`results/capacity_certificate.json`	The winning main model, portfolio, bindings, period ledger, and stated scope
`results/main.json`	Every main portfolio, terminal score, trace, and diagnostic search count
`results/controls.json`	Ten changed-assumption controls and their portfolio scores
`results/adaptive_controls.json`	Flexible-procurement results for three initial treasuries
`results/sweep.json`	Fifty-four finite configurations, 216 portfolio checks, and assertion counts
`results/stellar_ledger.json`	Constants, hypothetical inputs, calculated values, and recycling sensitivity
`CLAIMS.md`, `IMPLEMENTATION.md`	Claim boundaries and the proposed Alpha integration contract
`MANIFEST.sha256`	File-integrity checksums; neither a trusted signature nor a semantic proof

The main optimizer visits 72 cached states across its four portfolios, including terminal states. This is a diagnostic count, not a complete operation count or wall-clock performance measurement. It omits interpreter overhead, allocation, hashing, serialization, and the cost of supplying the source model. No computation-to-money conversion is inferred from it.

B Exact recurrences and independent comparison

B.1 Fixed-portfolio recursion

For a chosen portfolio with fee K, flow requirements (a, b), and delay d, define the terminal score

$$V_H(k, c, m) = (k, c). \tag{21}$$

For $t < H$, let

$$\mathcal{U}_t(k, c, m) = \{0, \ldots, \min\{\lfloor k/a \rfloor, \lfloor k/b \rfloor, c, m\}\}, \tag{22}$$

except that $\mathcal{U}_t = \{0\}$ for $t < d$. Then

$$V_t(k, c, m) = \max_{u \in \mathcal{U}_t}^{\text{lex}} V_{t+1}(k + u, c - u + r, m - u). \tag{23}$$

The initial state is $(k_0, c_0 - K, m_0)$. The program stores maximizing decisions and constructs the reported trace by executing them. Finiteness of the horizon, stock bounds, and action sets makes the recursion exact in this model. No approximation or trained value function is used.

The separate comparator instead propagates reachable triples forward, explicitly trying all bounded purchases and rejecting those that violate the flow or timing conditions. Its maximum terminal pair is compared with the recursive score. Agreement is evidence against some implementation defects; it does not make their common physical assumptions true.

B.2 Flexible procurement

The flexible state additionally contains the activation dates (d_E, d_F), with a distinguished value for an unpurchased project. At each period, the planner chooses any affordable subset of not-yet-purchased, permitted projects. A project purchased at t becomes usable at $t + d$, including the same period when the declared delay is zero. The planner then chooses construction subject to the currently active recipe, material, and cash *after* purchase. The old recipe remains available while a new project is pending.

For the three-unit treasury, an optimal path has end-of-period capacities

$$(1, 1, 2, 4, 8, 10) \tag{24}$$

and end-of-period cash

$$(5, 3, 4, 4, 2, 2). \tag{25}$$

Both projects are purchased in period one, after the first receipt has settled. Their fee is four; nine new modules cost nine; six receipts contribute twelve. Thus $3 + 12 - 4 - 9 = 2$ cash remains. This path changes a timing restriction, not the truth of any receipt, permission, or flow constraint.

B.3 What the finite optimum does not optimize

Both source formulations prohibit partial multi-period construction and purchases of finished modules from an external market. The exact flexible optimum is unrestricted *within its declared finite action model*; it is not the optimum over every imaginable industrial or financial arrangement. The prescribed project catalogue

contains only two nonincreasing requirement changes. A natural discovery study must disclose this supplied engineering and investigate alternatives beyond it where authorized.

The units of energy service, fabrication service, cash, material, and period are intentionally abstract. Mapping a module to actual hardware requires measured inputs, capacity, reliability, lead time, depreciation, maintenance, and a bill of materials. No rescaling of a dimensionless synthetic result supplies those measurements.

C A claim–evidence register

Claim	Evidence	Boundary
A binding resource cut cannot be bypassed with cash alone	Proposition 3.1; conservation inequality	Only while resource requirements and availability remain as stated
Either single project cannot bootstrap expansion from one module	Proposition 4.1; exact factorial	Indivisible one-period construction; no external finished-module purchase
The joint change reaches nineteen modules	Executed six-period trace; separate checker and forward optimum	Supplied project effects, demand, receipts, and resource model; not a real return
All six modeled service deliveries remain present	Period-level outcome and accounting checks	Service reservation and its net costs are stipulated, not independently measured
Three initial cash units can reach ten modules with flexible timing	Separate recursive and forward exact planners; retained purchase trace	End-of-period modeled receipts; no debt, uncertainty, or natural financial calibration
The reported sweep agrees with a separate method	54 configurations, 216 portfolios, 12,528 semantic checks	Finite regression coverage, not independent discovery campaigns
A repeatable scalable schedule supports conditional expansion	Theorem 5.1 and induction	Only within the stated scaling domain, supply, maintenance, permission, and bootstrap assumptions
High private value requires useful realization and capture	Separate enterprise accounting and a proposed observed-conversion study	No estimated revenue, asset valuation, commercial probability, or expected return
Stellar-scale power entails large material and thermal commitments	Dimensioned calculations with IAU/CODATA inputs	Hypothetical efficiency, temperature, geometry, and areal mass; not a validated design
Recycling does not create net stock	Explicit geometric gross-throughput model	Its gross processing count is not installed mass or new feed
AGI Alpha Agent is a concrete implementation target	Pinned repository guide and a proposed versioned adapter contract	Current preparation/execution scope is not industrial autonomy or delivered new integration
SUCCESSOR Ω can improve the discovery of productive paths	H1–H6, matched complete campaigns, staged milestones	Research hypotheses; no frontier-model or real-world superiority result

D A practical mission contract for Alpha and SUCCESSOR Ω

The first pilot should produce one useful outcome and one measured next capability, rather than implement a civilization-scale controller. The following is a proposed assignment, not a deployed service:

> Given an approved public computational workload and an independently checkable output contract, reproduce the incumbent, diagnose a costly restriction, construct an executable alternative, and test the claimed mechanism against strong same-permission controls. Count all construction and review costs. If a result is accepted, retain its supported capability and demonstrate one additional authorized outcome using the actually released or acquired resources. Preserve every previous obligation.

Admission freezes the beneficiary's outcome, required accuracy and timing, source rights, available tools, permissions, budget, and a comparison horizon. Development retains unsuccessful alternatives and may propose representations, experiments, algorithms, or changed work decomposition. Confirmation freezes a finalist and uses fresh conditions. Operational admission is a separate decision tied to the exact artifact and current context.

A useful result report distinguishes *proposed*, *constructed*, *executed*, *independently checked*, *accepted*, *paid*, and *commissioned*. Not every mission needs every state. A free numerical result can be valuable without a payment; a simulated receipt cannot be promoted to real cash; a drafted purchase cannot be promoted to capacity. A claim about full economic conversion requires all relevant observed transitions.

A capability-to-capacity record should retain the parent mission, source and model versions, the changed restriction, the original-outcome bridge, resource and cash ledgers, recipient capacity, commissioning test, and conditions requiring retirement or requalification. Use additive versioned schemas. Historical imports remain inert, original papers remain unchanged, and old zero-allowance requests cannot acquire positive spending authority through reinterpretation.

The immediate research implementation is digital and bounded. Physical experiments require separately qualified domain adapters, partners, operating approvals, and independently measured outcomes. Neither the proposed schema nor the supplied simulator authorizes real financial execution, broad physical control, or off-world construction.

References

Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025. https://doi.org/10.1093/epolic/eiae042. Open working version: https://economics.mit.edu/sites/default/files/2024-04/The%20Simple%20Macroeconomics%20of%20AI.pdf.

Philippe Aghion, Benjamin F. Jones, and Charles I. Jones. *Artificial Intelligence and Economic Growth*. NBER Working Paper 23928, 2017. https://doi.org/10.3386/w23928. Open text: https://www.nber.org/system/files/working_papers/w23928/w23928.pdf.

Margaret A. Boden. Creativity and artificial intelligence. *Artificial Intelligence*, 103(1–2):347–356, 1998. https://doi.org/10.1016/S0004-3702(98)00055-1.

Vincent Boucher. *Transformational Intelligence: Learning to Change the Game*. MONTREAL.AI research manuscript, October 2026. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/main/docs/Transformational_Intelligence.pdf.

Vincent Boucher. *Beyond the Original Problem: Transformational Intelligence for Expanding What Is Possible*. MONTREAL.AI research manuscript, October 19, 2026. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/main/docs/Beyond_the_Original_Problem.pdf.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372, 2021. https://doi.org/10.1257/mac.20180386.

Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. *Preserving Statistical Validity in Adaptive Data Analysis*. arXiv:1411.2664, 2014. https://arxiv.org/abs/1411.2664.

Kevin Ellis et al. *DreamCoder: Growing generalizable, interpretable knowledge with wake-sleep Bayesian program learning*. arXiv:2006.08381, 2020. https://arxiv.org/abs/2006.08381.

Google. *OR-Tools: Constraint Optimization*. Official documentation, accessed October 19, 2026. https://developers.google.com/optimization/cp.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. Open version: https://arxiv.org/abs/1810.08240.

Shengran Hu, Cong Lu, and Jeff Clune. *Automated Design of Agentic Systems*. arXiv:2408.08435, 2024. https://arxiv.org/abs/2408.08435.

International Astronomical Union. *Resolution B2: Re-definition of the astronomical unit of length*. 2012. Resolution text hosted by Observatoire de Paris: https://syrte.obspm.fr/IAU_resolutions/IAUResol_2012_0.html.

International Energy Agency. *Energy and AI*. Paris, 2025. https://www.iea.org/reports/energy-and-ai.

N. S. Kardashev. Transmission of Information by Extraterrestrial Civilizations. *Soviet Astronomy*, 8:217–221, 1964. https://adsabs.harvard.edu/pdf/1964SvA.....8..217K.

Philip T. Metzger, Anthony Muscatello, Robert P. Mueller, and James Mantovani. Affordable, Rapid Bootstrapping of the Space Industry and Solar System Civilization. *Journal of Aerospace Engineering*, 26(1):18–29, 2013. https://doi.org/10.1061/(ASCE)AS.1943-5525.0000236. Open manuscript: https://arxiv.org/abs/1612.03238.

MONTREAL.AI. *AGI Alpha Agent: Find your Alpha. Build your SUCCESSOR Ω.* Repository README, commit `9c7484d05cd257c657276d52f274e085488e9dcc`; inspected October 19, 2026. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/9c7484d05cd257c657276d52f274e085488e9dcc/README.md.

MONTREAL.AI. *Useful Alpha: a capability worth keeping*. Pinned flagship guide, same commit and inspection date. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/9c7484d05cd257c657276d52f274e085488e9dcc/docs/agent/SUCCESSOR_ALPHA.md.

National Institute of Standards and Technology. *2022 CODATA Recommended Values of the Fundamental Physical Constants*. Stefan–Boltzmann constant; accessed October 19, 2026. https://physics.nist.gov/cuu/Constants/Table/allascii.txt.

Alexander Novikov et al. *AlphaEvolve: A coding agent for scientific and algorithmic discovery*. arXiv:2506.13131, 2025. https://arxiv.org/abs/2506.13131.

Andrej Prša et al. *Nominal Values for Selected Solar and Planetary Quantities: IAU 2015 Resolution B3*. arXiv:1605.09788, 2016. https://arxiv.org/abs/1605.09788.

Graeme Ritchie. The transformational creativity hypothesis. *New Generation Computing*, 24(3):241–266, 2006. https://doi.org/10.1007/BF03037334. Author manuscript: https://homepages.abdn.ac.uk/g.ritchie/pages/papers/newgencomp.pdf.

Paul M. Romer. Endogenous Technological Change. *Journal of Political Economy*, 98(5, Part 2):S71–S102, 1990. https://doi.org/10.1086/261725. Open text: https://web.stanford.edu/~klenow/Romer_1990.pdf.

Nathan J. Szymanski et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature*, 624:86–91, 2023. https://doi.org/10.1038/s41586-023-06734-w. Author manuscript: https://perssongroup.lbl.gov/papers/szymanski_autonomous_laboratory_2023.pdf.

J. von Neumann. A Model of General Economic Equilibrium. *The Review of Economic Studies*, 13(1):1–9, 1945. https://doi.org/10.2307/2296111. Open text: https://piketty.pse.ens.fr/files/VonNeumann1945.pdf.

Jason T. Wright. *Application of the Thermodynamics of Radiation to Dyson Spheres as Work Extractors and Computational Engines, and their Observational Consequences*. arXiv:2309.06564, 2023. https://arxiv.org/abs/2309.06564.

Tianbao Xie et al. *OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments*. arXiv:2404.07972, 2024. https://arxiv.org/abs/2404.07972.

Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. *Darwin Gödel Machine: Open-Ended Evolution of Self-Improving Agents*. arXiv:2505.22954, 2025. https://arxiv.org/abs/2505.22954.

Jenny Zhang, Bingchen Zhao, Wannan Yang, Jakob Foerster, Jeff Clune, Minqi Jiang, Sam Devlin, and Tatiana Shavrina. *Hyperagents*. arXiv:2603.19461, 2026. https://arxiv.org/abs/2603.19461.