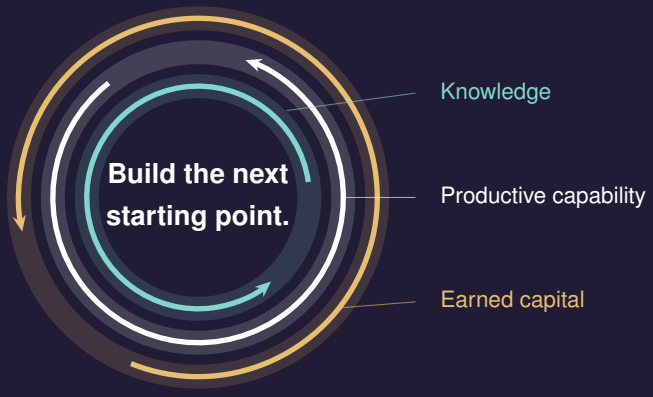
AGI ALPHA

FLYWHEEL

Inventing What Becomes Possible

TRANSFORMATIONAL INTELLIGENCE ∞



The output of one invention can become the means of the next.

Vincent Boucher

MONTREAL.AI • Montréal, Québec, Canada

2 November 2026 • Research edition 1.0

THE THESIS

Intelligence can build the conditions for its next useful invention.

Abstract

An autonomous enterprise can complete more work without becoming a better inventor, and can invent a useful tool without creating a viable business. AGI ALPHA FLYWHEEL studies the connection: how accepted work can leave behind reusable knowledge, qualified productive means, and available earnings that make a more consequential successor attainable. Building on *Transformational Intelligence*, we make the complete productive arrangement the object of search: the outcome, its production and verification, participation, financing, and allocation of obligations. The central test is whether a gain survives into subsequent work and changes what can be achieved after its full burden is counted.

We develop an architecture for this succession, a state model that separates capability from evidence and cash from profit, and three distinct lineage claims: a successor may be financed by prior earnings, enabled by an inherited instrument, or discovered more effectively by an improved search procedure. A reproducible finite construction shows how retained means and collected earnings can jointly cross a funding threshold even when neither suffices alone. A second experimental design separates better equipment from better invention. Public repository evidence supplies working constructions, conventional-method ties, information limits, and negative economic results; it does not establish an autonomous new industry or the US$3 billion trailing-year profit objective. The contribution is a synthesis, a precise evaluation discipline, and a research program for turning invention into a repeatable source of new productive possibilities.

The question

What valuable activity becomes economically possible when intelligence can construct the missing means of producing, verifying, and organizing it?

Read in three passes

For the idea, read Sections 1–4 and Figure 1. For the formal and economic model, read Sections 5–7. For evidence, comparisons, and the research frontier, read Sections 8–11.

Contribution and status

This is an architecture and research synthesis with an exact illustrative model and a source-bound evidence review. The illustrative arithmetic is executed by the companion program. Previously published experiments are cited in their original scopes; this paper does not recast them as new confirmation. The public source snapshot is `e67aa0903158`. Prospective capital is a design variable; US$100,000 remains a historical reference, and US$3 billion is an unachieved profit milestone, not a promised result or a ceiling.

Keywords transformational intelligence; productive arrangements; invention; capability inheritance; autonomous enterprise; open-ended discovery; funded succession.

1 Build the next starting point

Imagine asking a question about the past that nobody planned to measure: how did a particular production condition evolve before a failure? A simple archive lookup returns whatever the archive already records. A more ambitious response asks whether independent traces, a measurement model, and a calibrated reconstruction can create a defensible new series. Success would produce a useful answer and an instrument for answering a class of questions that had previously been uneconomic to investigate. Failure might reveal that the required distinction is absent from the available evidence. More fluent inference would not repair that absence.

Retrospective inference itself is established practice; the research question here concerns autonomous construction and qualification of a useful measurement process. This is a candidate direction, not a claim that arbitrary historical facts can be recovered. It illustrates the central object of AGI ALPHA: *constructing a new route to an important outcome*. The route may depend on a measurement, a representation, a production process, a coordination rule, an agreement, or a combination of them. The opportunity and the means of serving it are designed together.

Transformational Intelligence defines the relevant shift: identify a restriction in the current problem formulation, construct a different setup, preserve the authorized purpose, and show that the change matters [1]. The present paper follows that change through its consequences. Does it leave a reusable means? Can that means be recognized and used in a different task? Can useful delivery cover its construction and continuing obligations? Does the resulting capability or available surplus enable the next opportunity? Does the procedure for finding such opportunities improve?

A flywheel turns when a gain survives into the next decision.

The word *flywheel* names a proposed causal structure. It is not evidence of momentum. A growing archive may become harder to search. An expensive instrument may never be reused. A profitable delivery may leave its cash unavailable. A successor may succeed because a specialist supplied the answer. Each possibility changes what the enterprise should retain and what the researcher may claim.

1.1 Three things worth accumulating

Knowledge includes supported relations, applicability conditions, counterexamples, and procedures for choosing useful next experiments. Its value depends on whether it changes a decision or reduces justified work. **Productive capability** consists of usable instruments, interfaces, processes, or installed assets that can deliver accepted outcomes in a specified environment. **Earned capital** consists of resources actually available from settled useful delivery, after restrictions and obligations. These stocks can complement one another, but none is a substitute for the others.

The most valuable output of an intelligent process may be a durable environment in which a difficult inference is no longer repeatedly needed. A source-preserving production process can make later claims cheaper to inspect. A calibrated fixture can make later experiments interpretable. A useful interface can make previously incompatible components composable. These are candidates for reducing the cost of future intelligence through construction, consistent with the broader research on reusable discovery and productive capacity [2, 3].

1.2 A consequential objective with testable steps

The research objective is a sustained expansion of valuable outcomes achievable through voluntary cooperation. The repository's commercial milestone is US$3 billion of profit after complete costs over an observed trailing 365-day window [4]. The ambition directs opportunity selection; it supplies no evidence that a particular opportunity is large or viable. A small, decisive experiment can be appropriate when it tests a constraint on a consequential activity.

Ordinary profitable delivery remains worthwhile even when capability is fixed. Conversely, an early instrument may deserve finite research investment before it earns standalone revenue. The enterprise

chooses among delivery, acquisition, research, reuse, construction, and holding resources. It need not force an invention into every transaction.

2 Invent the complete productive arrangement

Definition 1 (Productive arrangement)**.** *A productive arrangement specifies an intended beneficiary outcome, the means of producing and accepting it, the participants and their outside options, the terms of access and exchange, the financing schedule, and the allocation of remaining work and risk.*

This definition makes commercial design part of construction. A technically superior output can fail because verification is unaffordable, the beneficiary cannot integrate it, the producer must fund an excessive delay, or another provider can supply the same accepted result at better terms. An agreement can fail because the promised means do not exist. Joint search can change these interdependencies [5].

The unit of comparison is therefore a complete route to the same accepted outcome. The incumbent may use humans, conventional software, frontier agents, partners, different financing, or a process that avoids the proposed invention altogether. It must receive comparable material permissions. A restricted baseline can diagnose the effect of a particular restriction; it cannot establish superiority to competent unrestricted practice.

2.1 Purpose, formulation, and evaluation

Three objects govern a candidate. The **purpose** states the authorized outcome and material requirements. The **formulation** states the current representations, information procedures, dependencies, and means. The **evaluation instrument** supplies evidence about whether the outcome was achieved. A transformation may revise the formulation. If an evaluation instrument proves inadequate, it may be improved under an explicit bridge between old and new measures; the revision must not retroactively erase a failure or silently replace the purpose [1, 6].

For example, a request for absolute local stress is not satisfied by recovering a relative average more accurately. The relative average may be useful for a different authorized purpose, but that is a separately specified offer. Likewise, improving a method of checking a claim does not license weaker acceptance. These distinctions protect the meaning of an outcome while leaving its means open to invention.

2.2 The missing means may be a connection

Represent a candidate instrument by

$$m : (\text{inputs}, \text{state}, \text{resources}) \longmapsto (\text{output}, \text{new state}, \text{evidence}). \tag{1}$$

Its contract includes prerequisites, units, input and output semantics, authority, timing, costs, failure conditions, and invalidation rules. Two qualified instruments do not automatically form a qualified composition. Their clocks may disagree, their quantities may refer to different objects, or both may consume a resource that was counted only once. The connector that resolves such an incompatibility can be the enabling invention.

This calls for a graph with *joint prerequisites*: several parents can be necessary together. A measurement and a calibration rule may jointly enable an estimate; an estimate and an acceptance agreement may jointly enable delivery. Individual components can have little value before the bundle works. The system should therefore compare bounded compositions, including ones whose individual edits are locally unhelpful, while charging the search and integration costs.

2.3 A compact arrangement record

Each serious candidate needs a short, reviewable record: the beneficiary outcome; the strongest complete alternative; the suspected binding constraint; the executable construction; independent acceptance; partici-

pants, rights and refusal options; lifecycle costs and payment timing; and a named successor that could use the resulting asset. The record's function is to expose a consequential next experiment. It is not a requirement to grow an administrative framework before doing useful work.

3 The flywheel architecture

Figure 1 makes the causal proposal visible. Selection starts with consequential outcomes, finds a binding condition, and constructs or acquires means that could change it. Independent evidence determines what actually works. Delivery on viable terms can produce both reusable capability and settled surplus. Retention changes the starting state for subsequent work. Discovery itself is evaluated separately.

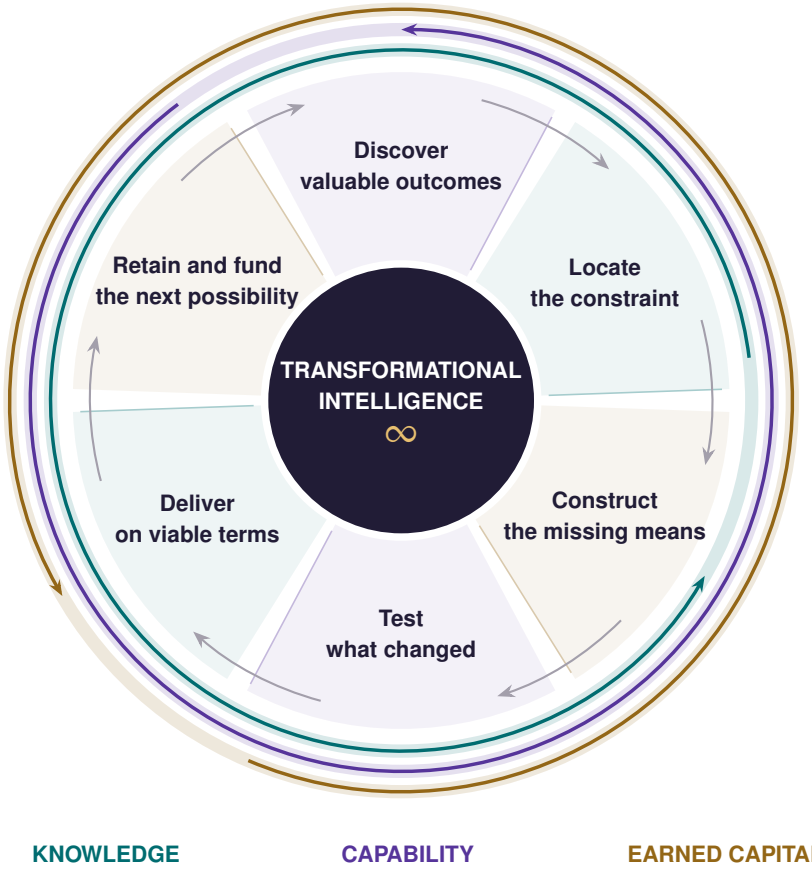


Figure 1. Three stocks, one testable succession. Knowledge improves selection; qualified means change what can be done; available earnings can finance the next action. The outer responsibilities may be revisited, combined, skipped, or performed in parallel. Arrows express hypotheses to test, not automatic growth. Every retained asset carries a scope; every new delivery requires current acceptance; every action requires current authority.

The architecture has three feedback paths. The **production path** improves a route to accepted output. The **economic path** selects viable exchanges and finances further work from permitted resources. The **discovery path** revises how useful arrangements are found and constructed. One path may improve while another stalls. Keeping them distinct permits a profitable service to continue while an unhelpful research method is retired.

3.1 Inheritance with a boundary

A successor may inherit an executable instrument, evidence of its earlier behavior, an applicability model, and a constructor of related instruments. These are different assets. Historical evidence can

support reusable generic checks, but the new mission still needs checks against its actual inputs, outputs, environment, and requirements. Permission is evaluated for the current action; it is not inherited from the existence of a tool [7].

Reuse should invalidate when relevant code, dependencies, measurement semantics, environment, or acceptance changes. The correct response can be a targeted recheck rather than a full replay, provided the dependency model is itself adequate. An archive that cannot explain why an item is applicable cannot yet support a strong inheritance claim.

For the prospective historical sensor, the *instrument* would turn specified traces into a calibrated estimate with uncertainty. Its *applicability model* would identify which clock alignment, channel semantics, and observation conditions permit that estimate. A *constructor* would adapt the measurement process and its discriminating checks to a different target. A better *discovery procedure* would find such useful constructions more effectively from comparable starting means. One successful reconstruction establishes none of the later claims by itself; each requires its own target and comparison.

3.2 Retention includes deletion

The enterprise retains compact, executable value: what worked, where it applies, what failed, and what would change the decision. Redundant traces, obsolete adapters, and misleading associations impose retrieval and maintenance costs. Removing them can improve the system. Failed candidates remain in denominators and accounts even when bulky intermediate files are replaced by a reproducible failure witness and a short lesson. Compression is compatible with honest accounting; erasing incurred costs is not.

4 Mechanisms that change the starting point

The OUT research collection offers complementary ways to construct missing means. Table 1 is a selection guide, not a mandatory sequence. A competent system chooses the mechanism that changes the current constraint, including a strong established method with no claim of novelty.

Table 1. From a binding condition to a useful construction. Each mechanism retains its source's assumptions and validation limits.

Mechanism	What it can change	Decisive question
OUT-LEARN [8]	Construct observations or calibration that separate relevant possibilities.	Does an independent reference support the new information?
OUT-THINK [9]	Build a representation that preserves required answers and updates with less work.	Which present and future obligations survive compression?
OUT-DESIGN [10]	Construct a replacement or transition that makes needed capacity usable.	Can the system actually reach and operate the replacement?
OUT-STRATEGISE [11]	Choose information and construction together, or act across remaining uncertainty.	Does the complete plan work under the declared possibilities?
OUT-EXECUTE [12]	Reuse compatible boundaries while recomputing changed parts.	Are the interfaces sufficient and current outputs accepted?
OUT-EVOLVE [13]	Build means that help construct a different successor capability.	What loses feasibility or value when the inheritance is removed?
OUT-REFLECT [14]	Construct evidence with inspectable corruption and correction assumptions.	What shared error or violated fault bound defeats the assurance?

4.1 Transfer the relation, then test the bridge

Structural analogy maps relations rather than merely matching surface attributes [15]. A candidate transfer should identify the source mechanism, the corresponding target quantities, the required invariants, and a falsifying case. If a correspondence fails, the system can reject the transfer, construct a connector that establishes the missing condition, or redesign the route so that the condition is unnecessary. An analogy proposes a search direction; the target's acceptance evidence decides whether it works.

4.2 Turn a verified objection into a construction

A useful critic identifies a source-grounded omission, a contradiction, a boundary case, or a discriminating experiment. Another process adjudicates that objection before repair. A confirmed defect can require more than revising the answer: it can identify an unavailable observation, a defective evaluator, or an unnecessary dependency. The constructor then searches over changes that address the established cause.

Agent agreement is not independence of evidence. Several agents can share the same false source or missing observation. Their agreement may help inspect a derivation, but it does not manufacture external information. When observations cannot identify the requested quantity, the appropriate action is to improve observation, narrow an authorized claim, change production so the ambiguity ceases to matter, or stop.

Proposition 1 (An information boundary)**.** *Let two admissible worlds produce identical complete available observations but require different correct outputs. Any deterministic procedure using only those observations returns the same output in both worlds and therefore cannot be correct in both. Under an equal prior on two worlds with distinct exact answers, a randomized procedure whose randomness carries no additional world-dependent information has success probability at most one half.*

Proof. Identical observations induce the same output distribution. The probabilities assigned to the two distinct correct outputs sum to at most one. Averaging over the two worlds gives the bound. The deterministic statement is the corresponding special case. □

This elementary observation is useful because it directs invention toward the information channel rather than toward unsupported confidence. It is conditional on the stated worlds and observations; finding no affordable test does not establish that no test exists. The flagship paper's intervention example makes the same distinction and explicitly reports an unrestricted conventional planner that matches the constructed policy [1].

The hypothesis set is also open to challenge. Both initial explanations can be wrong; model misfit may require a third explanation or a different observation contract. Robustness over the current set remains conditional on its coverage of the relevant possibilities.

4.3 Produce evidence as part of producing the result

Dependencies, versions, assumptions, and acceptance conditions can travel with the output. This may reduce later reconstruction and make targeted revalidation economical. The benefit depends on completeness and correct interpretation of those records. A source pointer is not a proof, and a self-consistent provenance graph can still omit a decisive external fact. Evidence-producing construction should therefore be tested against deliberate omissions and changed dependencies, as well as ordinary successful cases.

5 What becomes reachable

Let the enterprise state be

$$x_t = (K_t, M_t, D_t, C_t, O_t, A_t), \tag{2}$$

where K_t is retained knowledge, M_t usable productive means, D_t evidence and exposure history, C_t cash, O_t obligations and commitments, and A_t current authority. Keep the discovery procedure θ_t separate. A

richer state and a better procedure are different interventions.

Given charter I, acceptance instrument V, and a declared resource budget B, define the following set for deterministic executable policies in the stated model:

$$\begin{aligned}\mathcal{R}(x; B, I, V) = \{y : \exists \pi \in \Pi(I, A),\ y = \mathrm{Out}(x, \pi),\ V(y) = 1,\\ \mathrm{Cost}_{\mathrm{life}}(x, \pi) \preceq B,\ \pi \text{ is fundable along its trajectory from } x\}.\end{aligned} \tag{3}$$

Here $\mathrm{Out}(x, \pi)$ is the actual modeled output of executing π from x. A stochastic extension must specify its required success probability and resource-risk criterion over the outcome distribution; a favorable realization alone is insufficient. The budget may be a vector of money, time, energy, and other resources; the model need not invent a conversion rate between them. Lifecycle cost includes obtaining and maintaining the means, failures, and residual obligations. A plan that can be paid for only after an unavailable future receipt is outside the fundable set even if its terminal profit would be positive.

Equation (3) is an existential definition. It does not imply that the implemented search procedure can find every feasible plan. Three statements remain separate: an outcome is possible under a model; a procedure can discover it within budget; and an actual execution has achieved it. A win over selected baselines establishes a relative gain. Strict expansion beyond *all* old routes requires a valid lower bound or exhaustive argument over a specified feasible set.

5.1 Capability gain is not automatically frontier expansion

An additional tool can preserve an old route and add a useful new one. But acquiring, integrating, selecting, and maintaining the tool consumes resources. Therefore even set inclusion is conditional: all old feasible outcomes remain feasible only if old routes stay available and the added burden consumes no binding resource. Real systems can become less capable through accumulation. End-to-end measurement, rather than tool count, determines the practical result.

The useful empirical question is often narrower and more answerable: does the candidate deliver the specified outcome at a lower complete cost, within a tighter deadline, under a new accepted condition, or at a budget where the strongest tested alternative fails? Report these dimensions explicitly. A scalar reward is appropriate only when its tradeoffs are justified and fixed before the comparison.

5.2 Three lineage relations

Financed by. A source-bound record shows which previously available resources paid for the successor's commitments. Earnings-funded, seed-funded, and externally financed actions receive distinct labels. Financing proceeds are not operating revenue or earned profit.

Enabled by. Removing a retained asset changes the successor's accepted outcome, complete cost, or admissibility under the declared comparison. A strong alternative without that asset may acquire another instrument, construct its own, or avoid the supposed prerequisite. Mere use of the asset is insufficient evidence of necessity.

Discovered better by. A revised discovery procedure finds or constructs better arrangements across fresh campaigns from matched starting capabilities, information, permissions, and resource envelopes. The relevant unit is the full discovery campaign, including unsuccessful candidates. A favorable descendant alone does not identify this effect.

These relations can join into a causal lineage, including jointly enabling parent sets. Their benefits need not add: two instruments can be complements or substitutes. Attribution must preserve that interaction rather than credit each parent with the entire successor result.

6 Economics that can finance a next step

Let recognized operating revenue and any separately recorded recovery income over an interval be R, and the complete incurred expense under a declared accounting policy be X. Then $\Pi = R - X$; financing proceeds are excluded from R. Separately,

$$C_t = C_0 + \text{nonfinancing receipts}_t + \text{financing received}_t - \text{cash outflows}_t. \tag{4}$$

Nonfinancing receipts include settled customer collections and collected recoveries. Current available cash is $U_t = C_t - Q_t$, where Q_t is the *nonoverlapping* protected amount for custody, commitments, required buffers, and obligations that must be funded before further receipts. Exact category definitions belong to the account implementation. A restriction on cash is not automatically another expense, and the same obligation must not be subtracted twice.

6.1 Four separate decisions

Commercial viability asks whether accepted exchange can produce complete-cost profit under the stated conditions. **Investment value** asks whether an incremental use of resources is worthwhile against its alternatives. **Affordability** asks whether current resources can fund it while preserving obligations. **Successor benefit** asks whether the inherited means actually help the subsequent task. A business can pass the first while failing any of the others.

Positive profit does not imply affordable expansion. Revenue can be recognized before collection; reserves can restrict settled cash; an attractive investment can exceed the remaining balance. The correct response is to defer or redesign the investment and preserve the profit result. Conversely, positive cash financed by opening equity can coexist with losses. Counting either balance as proof of commercial success would hide the relevant constraint.

6.2 Complete cost and honest scope

The enterprise must account for search and failed attempts, acquisition, onboarding, production, inference and tools, verification, correction, integration, maintenance, support, payment losses, financing terms, and taxes or provisions where applicable. Actual human work is included when it occurs. A hybrid operation is labeled as such; normal buyer participation is separately disclosed. Access through an included subscription does not establish zero economic resource cost. Unmeasured costs remain unknown or enter an explicitly labeled sensitivity range.

Costs are recognized once under a consistent policy. A model that expenses development immediately must not charge the same amount again as depreciation. A model that capitalizes assets needs the corresponding depreciation, impairment, disposal, and cash treatment. Funding sources and outcome attribution remain distinct from either accounting choice. The existing native account is the operational authority; analytical views and pedagogical examples do not create another spendable balance.

6.3 Voluntary participation and strategic response

For a participant i, a modeled participation condition is

$$\mathbb{E}[u_i(\text{arrangement}) \mid \mathcal{I}_i] \geq \mathbb{E}[u_i(\text{outside option}) \mid \mathcal{I}_i]. \tag{5}$$

This is a design condition, not observed consent. Information, alternatives, payment timing, loss exposure, and the right to decline matter. A complete offer may need to change verification or risk allocation before any attractive price exists. Competitors and self-service can erode prices even when production is inexpensive.

The primary productivity benchmark compares competent frontier AI or agents with competent entirely human workers using conventional software on the same complete task. Human-plus-AI providers, autonomous rivals, self-service, and no purchase remain additional commercial alternatives. Human labor

cost is not a guaranteed demand curve. Price, scope, acceptance probability, payment terms, capacity, and capital should be selected jointly rather than inferred from a fixed wage discount.

Claims of game-theoretic optimality must identify the game, information, feasible strategies, mechanism, participant responses, and objective. The bilateral-trade impossibility results of Myerson and Satterthwaite [16] show why desirable efficiency and incentive properties cannot simply be assumed to coexist under private information without the stated qualifications. A finite modeled optimum is useful when its assumptions apply; it is not a universally optimal enterprise.

6.4 The scale objective is an observer

For a fixed endpoint T after a full operating year, define

$$\Pi_{365}(T) = R([T - 365\text{d}, T)) - X([T - 365\text{d}, T)). \tag{6}$$

Revenue, valuation, financing proceeds, cumulative multiyear profit, and annualized short runs are different quantities. Report initial capital, received financing, remaining obligations, accounting policy, and simulation or actual-commerce status alongside this observer. Prospective capital can be chosen jointly with the arrangement; historical fixed-capital experiments retain their original conditions [4]. Neither more funding nor a larger numerical target creates demand, an informative measurement, or an accepted product.

7 A finite construction of funded succession

The following example isolates a possible complementarity between retained means and earnings. Its integer units and opportunities are authored for explanation. They are not market estimates, native account records, or evidence that a system discovered the arrangement autonomously. All costs in this finite model are immediate; no financing, taxes, depreciation, hidden maintenance, or obligations exist unless a case explicitly introduces them.

A predecessor starts with 10 units, spends 8 on its entire accepted delivery including instrument construction, and collects 14. Its profit is 6 and its cash is 16. The next accepted outcome would collect 20 after delivery. The cheapest route in the declared finite alternative set requires 18 units without the instrument and 15 with it. An unrestricted conventional constructor is also allowed to build the same instrument during the predecessor; it can tie the result.

Proposition 2 (A complementary funding threshold)**.** *Fix the accepted successor, timing, permissions, and declared alternatives. Let c_0 and c_1 be the minimum immediate funding requirements without and with an inherited instrument. Let b be reference available cash and $e \geq 0$ collected unrestricted predecessor earnings. With no additional financing, define $S_{ij} = \mathbf{1}\{c_i \leq b + je\}$ for instrument access i and earnings access j. If*

$$b < c_1 \leq b + e < c_0, \tag{7}$$

then $S_{00} = S_{10} = S_{01} = 0$ and $S_{11} = 1$.

Proof. The first inequality excludes both methods at cash b. The last excludes the no-instrument method at cash $b + e$. The middle inequality admits the instrument method at that cash. These are funding decisions; delivery still requires independent acceptance. □

The combined arm can spend 15 and, conditional on acceptance and collection, finish with $16 - 15 + 20 = 21$. Cumulative profit is $(14 - 8) + (20 - 15) = 11$, so final cash equals initial cash plus profit. The instrument's original construction is already included in 8; it is not charged again. The binary interaction $S_{11} - S_{10} - S_{01} + S_{00}$ is one. This is one threshold witness, not a general law of supermodularity or exponential growth.

Table 2. The same predecessor history, four successor access conditions. Quarantining earnings changes availability, not recognized profit. All amounts are illustrative units.

Instrument	Earnings access	Available	Needed	Admission
Withheld	Withheld	10	18	Defer
Retained	Withheld	10	15	Defer
Withheld	Available	16	18	Defer
Retained	Available	16	15	Fundable

7.1 Failure cases give the example its meaning

Withhold the 6 earned units from successor use and keep them explicitly restricted: recorded predecessor profit remains 6. Delay the 14-unit collection and, under the example's recognition rule, profit can still be 6 while cash is only 2. Reserve 2 of the settled 16 and only 14 remains available, again deferring the 15-unit action. If the funded successor fails acceptance after spending 15, it collects nothing; cash is 1 and total profit is -9. Failed work does not disappear.

If a conventional alternative needs only 15 without the inherited instrument, the alleged necessity of that instrument vanishes. If the conventional constructor creates the same instrument at the same predecessor cost, it matches the entire successful sequence. Both controls belong in the interpretation. The example demonstrates how capability and liquidity can interact; it establishes no unique advantage of AGI ALPHA and no improvement in the inventor.

The companion `examples/succession.py` evaluates these cases with exact arithmetic and checks cash conservation, preserved profit under cash restriction, and the conventional tie. It is a pure explanatory calculation and opens no business accounts. The four access treatments isolate a mechanism after a common predecessor. A separate whole-campaign comparison from equal initial resources is required to establish that discovering this arrangement was better than the strongest alternative.

8 How to test an improving inventor

An impressive output does not identify the cause of improvement. Table 3 separates claims that are often compressed into a single success label. These are evaluation categories, not a mandatory ladder for every useful transaction.

Table 3. The claim determines the experiment.

Claim	Required comparison or observation	What does not establish it
Useful construction	Current independent acceptance in its intended environment.	Code existence, an attractive diagram, or self-approval.
Commercial success	Accepted exchange, collections, complete incurred costs, and covered obligations.	Funding, revenue alone, or forecast demand.
Retained successor benefit	Fresh successor with and without inheritance; strong equal-content alternatives.	Reuse frequency or a later successful timestamp.
Better discovery	Fresh campaigns crossing discovery procedure and starting capabilities.	A better-equipped descendant alone.
Physical productive expansion	Installed, accepted output or capacity with lifecycle evidence.	A simulation or an engineering drawing.

8.1 Cross the inventor with its equipment

Let $Q(\theta, M)$ be a preregistered campaign outcome, or a vector of accepted outcomes and complete costs, for discovery procedure θ and initial means M. Compare all four cells:

	Original means M_0	**Retained means** M_1
Original procedure θ_0	$Q(\theta_0, M_0)$	$Q(\theta_0, M_1)$
Revised procedure θ_1	$Q(\theta_1, M_0)$	$Q(\theta_1, M_1)$

Improvement in $Q(\theta_0, M_1)$ isolates the value of richer starting means under a fixed procedure. Improvement in $Q(\theta_1, M_0)$ tests a revised discovery process from the same means. The combined cell can reveal interaction. If a justified scalar Q is used, its factorial interaction is

$$Q(\theta_1, M_1) - Q(\theta_1, M_0) - Q(\theta_0, M_1) + Q(\theta_0, M_0). \tag{8}$$

An interaction is not automatically positive. Keep acquisition costs, exposure histories, and the distinction between deployed assets and methods for discovering assets explicit. This crossed design is an evaluation proposal here, not a reported new experiment.

8.2 Three comparison boundaries

Purpose. Hold the complete outcome and material acceptance conditions common. If the purpose changes, register the new comparison. **Means.** Match initial resources, access, information, financing menus, and authority, and allow competent alternatives to construct different means. **Evidence.** Keep confirmation separate from exposed development data, include unsuccessful registered candidates, and evaluate the complete discovery burden.

Equal-content ordinary reuse is essential: the proposed retained method should not win solely because the control never received the same useful information. A direct specialist is essential when the task has strong known structure. An unrestricted complete alternative is essential when the claim concerns transformation rather than escape from an artificially restricted baseline. These comparisons can be expensive; the strength of a claim should match the comparison actually conducted.

8.3 Finite search and interpretable stopping

Before a confirmation campaign, fix its candidate population or generation rules, resource envelope, independent outcome checks, comparison arms, and stopping conditions. Include search and verification cost. Preserve the result when a budget is exhausted, prerequisites fail, or the incumbent wins. A subsequent campaign can revise the construction, but its new cases and resources are explicit. An endless search stopped at the first favorable result supplies neither a calibrated success rate nor evidence of reliable discovery.

Repeated tasks within one campaign can share causes and are not automatically independent samples. When uncertainty estimates are justified, use the campaign or another defensible unit of independence. Report exposure and selection. A model improved using a benchmark cannot treat its repeated exposure as fresh evidence of transfer.

8.4 An operating cycle that can choose not to invent

The practical cycle is to identify an important accepted outcome, compare complete routes, locate the binding condition, execute a finite discriminating construction, and evaluate both delivery and its costs. Collect only under the applicable acceptance and settlement rules. Preserve obligations. Retain useful instruments and short applicability lessons. Choose the next action from the actual updated state, including ordinary delivery, a different business, a new experiment, or holding resources.

This cycle is a policy interface, not an assertion that the entire research program is implemented. Repository components provide bounded mapping, construction, criticism, reuse, and native simulated accounting; their call paths and current limitations are documented publicly [4]. A named module does not prove that an autonomous agent discovered the method embodied in it.

9 What the public evidence currently supports

The evidence review is pinned to repository revision `e67aa0903158`. The three studies below answer different questions and their denominators are not pooled. They are reported published records inspected for this synthesis, not new market experiments or independent reruns performed by this paper.

9.1 A working arrangement with a conventional tie

Experiment 063 constructs accepted synthetic flexible-load planning artifacts, independently checks them, and connects them to simulated commercial accounting [17]. Individual subthreshold resources satisfy 0 of 4 technical requests; both ordinary unrestricted HiGHS pooling and the AGI ALPHA construction satisfy 4 of 4. A separately delegated generalist also constructs an accepted fresh two-event successor. These are computational planning results, not physical device operation.

The central authored account collects US$250, incurs US$179.35 of modeled total expense including its declared income provision, and reports US$70.65 net profit. Copycat and no-purchase scenarios each lose US$128. In a supplier-cost and failed-delivery scenario, US$2.40 profit does not satisfy the specified complete commercial acceptance. Actual demand and complete host/model billing remain unmeasured.

Crucially, the inherited successor route costs US$18 against US$3 for the ordinary route, with the same accepted work. The extra US$15 investment is funded from available modeled earnings, but its realized successor benefit is absent. The appropriate conclusion is to retain useful base delivery while rejecting that incremental investment claim. This record directly separates *financed by* from *helped by*; it does not establish a new industry.

9.2 A measurement witness that redirects invention

The next-series manufacturing study constructs a local observation operator and an ambiguity witness within declared elastic and residual-eigenstrain models [18]. In three authored interior geometries, two admissible constructed stress fields have central σ_{xx} values of -100 and $+100$ MPa while their declared full-support uniform through-direction averages vanish. The test exposes information absent from that particular measurement operator; it is not an observed 200 MPa error in a physical specimen.

The source application remains on hold because the intended local absolute-stress use needs a qualified measurement contract, calibration, matching spatial kernels, and independent same-specimen truth. No raw-row accuracy confirmation or physical qualification occurred. A learned reconstruction cannot recover a distinction that the assumed observations erase. The retained asset is a scoped identifiability instrument and a discriminating next measurement requirement, not a licensed physical service or measured commercial advantage.

9.3 Capital alone did not repair the finite portfolio

The R2 public portfolio reports 110 separately funded authored native scenarios over 36 modeled months [19]. The best final trailing-365-day profit is US$0 from non-entry; the worst is negative US$36,000. Neither the US$1 billion nor the US$3 billion profit milestone is attained. Twenty-three accounts retain uncovered funding conditions; arithmetic acceptance does not convert their obligations into funded success.

The scope matters as much as the result. Each account has eight authored prospects and at most three scheduled purchases per prospect across three years, with a maximum registered price of US$40,000. This

imposes a loose US$960,000 gross-contract bound per account over the entire interval. The US$1 billion initial-equity arm leaves demand unchanged. Consequently the study cannot validate a route to billions in annual profit, and its failure is not a proof that all possible businesses fail. It shows that enlarging capital without changing the limiting arrangement does not create the desired scale.

9.4 What these records change

The evidence supports usable computational constructions, accounting distinctions, strong conventional comparators, and precise reasons to reject or defer proposed gains. It supplies no demonstrated open-ended improvement of discovery, no physical industrial expansion, and no actual commercial profit for these experiments. This is actionable knowledge: stop buying a tool that loses, obtain the measurement that would distinguish the target, and seek a different complete arrangement when the registered market cannot express the objective.

10 Relation to prior work

The research does not originate transformational creativity, structural analogy, reusable programs, open-ended search, or recursive self-improvement. Boden's account distinguishes combination, exploration, and transformation of conceptual spaces [20]. Gentner formalizes relational structure mapping [15]. DreamCoder jointly develops reusable program abstractions and search capability [21]. POET co-generates challenges and solutions and uses transferred solutions as stepping stones [22].

More recent systems directly investigate editable agents and improvement procedures. The Darwin Gödel Machine maintains and evaluates self-modified coding agents [23]; Hyperagents makes both task and meta agents editable and reports accumulating meta-level improvements [24]. These are close precedents for research on improving the inventor. AGI ALPHA studies how such improvement procedures could interact with accepted production, continuing obligations, and fundable succession. Those economic and physical outcomes require evidence beyond the reported agent benchmarks.

Within the AGI ALPHA series, the foundational substrate connects agents, jobs, validators, tools, memory, settlement, and governance [25]. *Transformational Intelligence* treats the formulation as an object of discovery [1]; *Invent the Economy* develops voluntary viable arrangements [5]; *The World as an Output* and *Beyond the Production Frontier* extend the means of discovery and production [2, 3]. The OUT collection supplies more specialized constructions with their own conditional results. This paper integrates those ideas into a single account of successive enterprise states and asks for separate evidence of funding, enablement, and improved discovery.

The contribution is therefore an operational synthesis and evaluation design: couple qualified inheritance to real resource constraints and complete productive arrangements, then follow the benefit through a successor. The elementary propositions and authored example clarify that synthesis. They are not presented as new mathematics, a new learning algorithm, or proof that this architecture dominates other organizations.

11 The frontier to pursue

The destination is an enterprise capable of opening important classes of useful work. The present services and experiments are candidates and instruments for learning; they need not determine its permanent industry. Three prospective directions make the distinction concrete.

11.1 Retrospective measurement as a service

A system might construct a new measurement series from existing traces for a question that was not anticipated when the traces were collected. A decisive first experiment would require a reconstruction

target, a declared observation model, and independently withheld truth. Nonidentifiability, selective archives, calibration drift, and rights to use the traces are explicit failure conditions. Commercial viability would further require a beneficiary who accepts the scoped uncertainty and a complete delivery price that covers acquisition, checking, and support. The result could be an instrument for answering a class of new questions; arbitrary recovery of the unobserved is not the proposition.

11.2 A foundry for qualified scientific instruments

A discovery system might build a measurement or experimental interface, use it to resolve one consequential uncertainty, and transfer its constructor to a distinct scientific task. The decisive chain contains independently reproduced measurements, a physically or otherwise appropriately validated instrument, and a fresh successor that loses a measurable benefit without it. Calibration, material variability, construction time, and failure costs can dominate. A simulation can select a design; it cannot certify its physical realization. The OUT-EVOLVE and reusable-discovery programs supply starting constructions and stated limits [2, 13].

11.3 New productive cooperation

Familiar resources can be unable to deliver an important complete service because their interfaces, timing, verification, or risk allocation do not fit. A new arrangement might connect them into useful capacity. The decisive experiment compares complete delivery against competent existing coordinators and alternative organizations, including voluntary refusal. The invention must remove a consequential obstacle rather than merely redraw a known service. Installed output, accepted contracts, and residual support obligations determine whether technical coordination becomes productive expansion.

These directions are research candidates, not forecasts of new industries. Selection should favor consequential missing means and the experiments most likely to change a decision, without invented success probabilities or market values. If the original obstruction proves removable by an existing method, use that method. If the full arrangement remains unattractive, change the arrangement or leave the opportunity.

11.4 What infinity means

The infinity mark names an open-ended research horizon: new purposes, instruments, and productive arrangements can enter the search. It does not assert infinite information, computation, memory, capital, or guaranteed improvement. A monotone best-so-far score $q \leq q_{\max}$ cannot accumulate net gains of a fixed $\varepsilon > 0$ more than $\lfloor (q_{\max} - q_0)/\varepsilon \rfloor$ times. Genuine long-run progress must therefore be described through new accepted capabilities and changing opportunity sets, not an impossible endless rise on one bounded benchmark.

The stronger hypothesis is that learned construction can reduce the burden of future useful construction often enough to sustain a sequence of increasingly consequential achievements. The counterhypothesis is that search, verification, maintenance, imitation, and resource constraints absorb the gains. Fresh campaign comparisons, including losses and declining returns, distinguish them. The flywheel may accelerate, stall, simplify, or be replaced.

12 Conclusion

AGI ALPHA FLYWHEEL asks intelligence to leave behind more than completed work. A useful project can create an instrument, a better understanding of when it applies, and available resources for a different project. A further advance occurs when the enterprise becomes better at discovering those transitions. Each connection can be specified, implemented, and challenged.

The central design choice is to search over the valuable outcome and the means of achieving it together. The central scientific discipline is to distinguish funding from causal enablement and better equipment

from a better inventor. The central economic discipline is to make accepted value, complete costs, available cash, and continuing obligations agree.

The current evidence supports pieces of that construction and rejects several stronger claims. The next consequential demonstration is a connected sequence in which a new productive arrangement earns its continuation, an inherited means materially helps a distinct successor, and independent comparison reveals exactly why. The US$3 billion objective remains ahead of the evidence. The research task is to build the means that could make such an outcome attainable and to discover, through execution, which proposed means deserve to survive.

Build the next starting point. Then test what it makes possible.

A Reproduce, inspect, and extend

The publication package contains editable LaTeX, a scholarly bibliography, the vector flywheel, a finite exact-arithmetic example, and a public source map. Build from `docs/agi-alpha-flywheel` with `make`. A TeX Live installation with pdfLaTeX, BibTeX, the recommended fonts, and the listed LaTeX packages is sufficient. The standalone source archive can also be uploaded to a LaTeX editor with `paper.tex` selected as its main file.

Run the explanatory construction with:

```
python3 examples/succession.py -check
python3 examples/succession.py
```

This needs only the Python standard library. It performs no model calls, network requests, external transactions, or native account mutations. Its output is an illustrative model result, not a business receipt. The existing native financial engine remains the authority for repository operating experiments.

For the reviewed experimental records, use the exact source links in the bibliography and the commands in each original report. Reproducing their software may require their pinned dependencies and frozen source. This paper's example does not reproduce those experiments, validate physical claims, or turn the public records into live commercial operations.

A.1 Minimal evidence record for a future claim

A future connected demonstration should bind its intended outcome, complete alternatives, source revisions, exposure history, finite protocol and stopping rule, acceptance evidence, complete incurred costs, cash restrictions and obligations, and the actual successor. Record separately whether each link was funded, causally enabling, or evidence of improved discovery. Include the outcome that would invalidate reuse and the result of the strongest conventional route. These fields make a result inspectable without requiring disclosure of private documents, customer data, or execution credentials.

A.2 A note on authorship and evidence

This manuscript presents Vincent Boucher's AGI ALPHA research program. Included AI assistance was used for synthesis, drafting, code preparation, and technical checking. The implemented illustrative example is authored, not autonomously discovered experimental evidence. Source-specific findings remain attributed to their public papers and records. Publication and reproducible arithmetic do not establish peer review, independent market validation, or completion of the broader program.

References

[1] Vincent Boucher. Transformational intelligence: Learning to change the game. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/transformational-intelligence-202610.pdf. Research manuscript.

[2] Vincent Boucher. The world as an output: Constructive sufficiency for reusable discovery. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/the-world-as-an-output-20261020.pdf. 20 October 2026. Research manuscript.

[3] Vincent Boucher. Beyond the production frontier: Transformational intelligence from verified discovery to stellar-scale civilization. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/beyond-the-production-frontier-20261019.pdf. 19 October 2026. Research manuscript.

[4] MONTREAL.AI. AGI ALPHA flywheel: Developer instructions, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/AGENTS.md. Operating direction at source snapshot e67aa09031583f976a39c8406ae86585b02cde76; consulted 2 November 2026.

[5] Vincent Boucher. Invent the economy: Transformational intelligence for creating viable markets. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/invent-the-economy-202610.pdf. Research manuscript.

[6] Vincent Boucher. Beyond the original problem: Transformational intelligence for expanding what is possible. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/beyond-the-original-problem-20261019.pdf. 19 October 2026. Revised research manuscript.

[7] MONTREAL.AI. Transformational autonomous enterprises: Mission-relative revalidation for capability inheritance across four execution systems. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/transformational-autonomous-enterprises-v1.pdf. 23 October 2026. Research edition 01; prepared by Writer 4; SUCCESSOR Omega harness.

[8] Vincent Boucher. Out-Learn: Change what you need to know. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-learn-20261021.pdf. 21 October 2026. Revised research manuscript.

[9] Vincent Boucher. Out-Think: Transformational intelligence for inventing the space of reasoning. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-think-20261021.pdf. 21 October 2026. Research manuscript.

[10] Vincent Boucher. Out-Design: Creating the freedom to transform. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-design-20261021.pdf. 21 October 2026. Revised research manuscript.

[11] Vincent Boucher. Out-Strategise: Transformational intelligence for constructing the conditions of success. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-strategise-20261020.pdf. 20 October 2026. Research manuscript.

[12] Vincent Boucher. Out-Execute: Constructing the conditions of completion. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-execute-20261021.pdf. 21 October 2026. Revised research manuscript.

[13] Vincent Boucher. Out-Evolve: Building the means of discovery and production. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-evolve-20261020.pdf. 20 October 2026. Revised research manuscript.

[14] Vincent Boucher. Out-Reflect: Transformational intelligence for constructing self-correcting evidence. Technical report, MONTREAL.AI, October 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/out-reflect-self-correcting-evidence-20261021.pdf. 21 October 2026. Substantively revised research edition.

[15] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983. doi: 10.1207/s15516709cog0702_3. URL https://doi.org/10.1207/s15516709cog0702_3.

[16] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983. doi: 10.1016/0022-0531(83)90048-0. URL https://doi.org/10.1016/0022-0531(83)90048-0.

[17] MONTREAL.AI. Constructing the productive arrangement: AGI ALPHA research edition 063, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/transformational063/REPORT.md. 1 November 2026. Repository report; source snapshot e67aa09031583f976a39c8406ae86585b02cde76.

[18] MONTREAL.AI. Next experiment series: Executed results and admission decision, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/experiments/next_series_20261102_365c3d6c/REPORT.md. 2 November 2026. N3 model-instrument and measurement-scope results; source snapshot e67aa09031583f976a39c8406ae86585b02cde76.

[19] MONTREAL.AI. AGI ALPHA: R2 complete finite comparison, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/experiments/public_portfolio_20261102/evidence/economics-r2/report/REPORT.md. 2 November 2026. Report of 110 separately funded authored scenarios; source snapshot e67aa09031583f976a39c8406ae86585b02cde76.

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

[21] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sablé-Meyer, Lucas Morales, Luke Hewitt, Luc Cary, 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*. ACM, 2021. doi: 10.1145/3453483.3454080. URL https://doi.org/10.1145/3453483.3454080.

[22] Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O. Stanley. Paired open-ended trailblazer (POET): Endlessly generating increasingly complex and diverse learning environments and their solutions, 2019. URL https://arxiv.org/abs/1901.01753v3. Version 3, 21 February 2019.

[23] Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin Gödel Machine: Open-ended evolution of self-improving agents, 2025. URL https://arxiv.org/abs/2505.22954v3. Version 3, revised 12 March 2026.

[24] Jenny Zhang, Bingchen Zhao, Wannan Yang, Jakob Foerster, Jeff Clune, Minqi Jiang, Sam Devlin, and Tatiana Shavrina. Hyperagents, 2026. URL https://arxiv.org/abs/2603.19461v1. Version 1, 19 March 2026.

[25] Vincent Boucher. AGI ALPHA: A scalable substrate for intelligence organizations. Technical report, MONTREAL.AI, April 2026. URL https://github.com/MontrealAI/AGI-Alpha-Invention-Flywheel-Lab/blob/e67aa09031583f976a39c8406ae86585b02cde76/docs/research/agi-alpha-foundation-20260430.pdf. 30 April 2026. Research manuscript.