Invent the Economy

Transformational Intelligence for Creating Viable Markets

Vincent Boucher
MONTREAL.AI
Montréal, Québec, Canada

October 2026 · Research manuscript

Abstract

An AI can improve a product while leaving intact the arrangement that makes the product uneconomic to create. We investigate a different object of invention: a technical capability and a voluntarily accepted commercial arrangement that make previously unviable work worth undertaking. This applies transformational intelligence to market formation itself, rather than treating the market, customer boundary, financing sequence, and unit of compensation as fixed. We formalize an economic-possibility package, separate technical correctness, participant benefit, profitability, liquidity, and realized exchange, and derive conditional thresholds for shared construction and native-service overhead. An executed exact reference model illustrates joint emergence: four hypothetical buyers cannot support separate developments, and one buyer cannot support a shared core; coordinated commitments together with the shared core yield an expected contribution of $120,000 under an assumed 80% qualification probability. A successful realization yields $300,000, but $900,000 of prefunding is required when delivery costs precede payment. A $40,000 per-delivery route burden reverses the expected gain. A separately implemented event-ledger checker agrees with the reference decisions across 2,240 parameter cases. A staged reuse study shows why positive expected contribution need not be financeable from earlier earned surplus. We also derive a conditional abundance paradox: greater technical productivity can reduce token-denominated transaction activity when revenue is tied to operations that the technology eliminates. The proposed AGI Alpha Agent / SUCCESSOR Ω architecture instead searches for new deliverable capabilities, supporting counterparties, and evidence-backed service relationships. These are analytic and supplied-model results, not autonomous business creation, real buyer commitments, or a token-price forecast. The central research target is intelligence that helps construct new economic possibilities—and can demonstrate both the value created and the limits of its capture.

Keywords: transformational intelligence; market formation; capability synthesis; complementarities; economic feasibility; $AGIALPHA; SUCCESSOR Ω.

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

1 The market is not always the input. Sometimes it is the invention.

Four organizations can each have a valuable problem, each rationally refuse to finance an expensive solution alone, and collectively leave a useful capability unbuilt. A better task executor does not necessarily change this outcome. Neither does a larger population of agents competing for the same fragmented requests. The missing construction can be simultaneously technical and institutional: a method that serves several requirements, a set of credible buyers, a division of costs and risks, and a delivery sequence that someone can actually finance.

The conventional instruction is: *find a customer, then solve the customer's problem*. The question developed here is: *which shared capability and voluntary arrangement would make a collection of currently unserved needs economically realizable?* An AI might discover that apparently different requests depend on a common mathematical construction; that one accepted artifact can support several adaptations; or that the route from funding to delivery is the restrictive assumption rather than the algorithm itself.

Transformational intelligence recognizes when the original problem formulation is unnecessarily restrictive—and discovers a fundamentally different game to play.

Here the game to reconsider is not only how to solve a task. It is what capability to build, who would benefit, and which arrangement makes its construction and delivery viable.

"Invent the economy" is deliberately an ambitious title with a bounded meaning. It does not mean inventing people, manufacturing their consent, or granting an AI authority to reorganize society. It means constructing new, purpose-grounded opportunities for voluntary production and exchange. A new market is not established merely because a model supplies a name, a business plan, or an attractive revenue spreadsheet.

1.1 A different unit of invention

Our prior papers moved from formulation discovery to persistent capabilities, production-frontier changes, and an improvement-for-capacity pilot (Boucher, 2026a,b,c,d). The latter asks whether a verified improvement can earn a useful resource for another mission. That remains one possible move. The present manuscript makes a broader object editable: the *economic possibility* in which a technical construction, its buyers, its resource requirements, and its contractual realization fit together.

This changes what a successful successor retains. A code artifact alone is insufficient. The reusable asset can include a verified method, conditions of validity, an independently implementable checker, permitted licenses, adaptation procedures, and experience about actual delivery costs. Those resources may change whether the next service is viable. Their value is conditional on transfer, maintenance, and a real beneficiary, not on the length of the agent's memory.

The proposal also addresses a mismatch between intelligence and monetization. A system that eliminates ninety-nine out of one hundred recurring operations improves productivity, but may reduce a business's revenue if its only product is billable operations. A native token attached mechanically to those operations can face the same conflict. A transformational system should not have to preserve unnecessary work to make its economic mechanism look successful.

1.2 Contribution, novelty, and evidence

Complementarity, shared fixed costs, advance commitments, platform coordination, and endogenous innovation have established literatures (Milgrom and Roberts, 1990; Kremer et al., 2020; Rochet and Tirole, 2003; Romer, 1990). We do not claim to have invented these mechanisms, a new general equilibrium theory, or a universally superior market designer. Likewise, program synthesis and automated agent design already make parts of the technical problem-solving process editable (Ellis et al., 2020; Hu et al., 2024; Novikov et al., 2025).

The contribution is an operational synthesis with four connected parts: a joint capability–arrangement

search problem; conditional viability and financing results that make structural attribution explicit; an executed exact-model laboratory with failure controls; and a proposed implementation and evaluation protocol linking this search to SUCCESSOR Ω and to a separately measured $AGIALPHA service economy. The missing capability must ultimately be constructed, the counterparties must actually agree, and a useful outcome must be delivered. None of those achievements is inferred from the reference arithmetic.

2 An economic possibility is an executable, mutually acceptable package

2.1 Separate purpose, means, parties, and evidence

Let a purpose contract be

$$\mathcal{I} = (\mathcal{O}, \mathcal{H}, \mathcal{A}, \mathbf{B}, \mathcal{Z}, H), \tag{1}$$

where $\mathcal{O}$ describes authorized outcomes, $\mathcal{H}$ nonnegotiable requirements, $\mathcal{A}$ permissions, $\mathbf{B}$ resources, $\mathcal{Z}$ exogenous conditions, and H the horizon. Different counterparties can have different purpose contracts. A coalition does not erase their individual requirements or convert a collective surplus into consent.

A proposed economic possibility is

$$\mathcal{P} = (T, C, A, F, R, W), \tag{2}$$

where T is a technical capability and its realizations, C a set of actual prospective counterparties, A a proposed arrangement of transfers and responsibilities, F a financing and delivery schedule, R a rights and authority record, and W evidence linking every claim to the appropriate outcome. The notation separates objects; it does not imply that the components can be chosen independently.

Definition 2.1 (Joint economic reformulation)**.** *Relative to a named incumbent, a joint economic reformulation changes a technical problem-solving setup and a material arrangement for making or delivering its outcome. A* demonstrated useful reformulation *additionally satisfies the relevant purposes, obtains the necessary real acceptance, and produces a defensible improvement after full resource and obligation accounting.*

A coalition that simply buys four copies of the same incumbent product can be useful without meeting this definition. A technically new method that has no accepted commercial arrangement can be valuable research without being a newly formed market. Classification of the change and evidence of its usefulness are separate questions.

2.2 Five gates, not one success score

Table 1: Evidence advances through distinct claims. A later label cannot be inferred from an earlier one.

Claim	Necessary evidence	Still not established
Technical possibility	Executable artifact and domain-specific outcome checks	Useful delivery economics
Modeled viability	Explicit costs, risks, participant constraints, and feasible cash schedule	Real demand or binding acceptance
Accepted arrangement	Authorized counterparties accept specific terms and risk	Successful implementation or payment
Realized exchange	Accepted delivery, actual receipts, liabilities and complete costs	Sustained market formation
Repeatable market	New independent adoption or renewal under viable terms	Universal scalability or token appreciation

A model output containing a buyer name and a willingness-to-pay estimate belongs to the first stages of investigation, not to the acceptance stage. Source-grounded public evidence can motivate a contact or experiment. Only the counterparty's authorized action establishes acceptance. The agent must never fabricate that action to close a promising plan.

2.3 Attainable economic outcomes, not a single valuation

For a formulation G that fixes what may be constructed and which arrangements may be considered, define

$$\mathcal{V}_I(G, \mathbf{B}) = \{\mathcal{P} : \mathcal{P} \text{ is realizable within } G, \mathbf{B} \text{ and meets the stated admissibility conditions}\}. \tag{3}$$

Admissibility includes physical feasibility, required outcomes, rights, cash timing, and individual acceptance at the evidential level claimed. A discovery process changes which packages can be found; a new capability or arrangement can change which packages are attainable at all under the stipulated formulation. These are not identical achievements.

A profitable baseline is an achieved reference, not a bound on every possible business. To establish a formulation-dependent expansion, one needs either a valid exclusion from the old attainable set or appropriately scoped comparative evidence. The exact model below supplies an exclusion only under its explicit information, construction, commitment, and financing rules. Natural applications will generally have weaker, empirical bounds.

3 When neither better technology nor better coordination is enough

3.1 A deliberately transparent shared-construction model

Consider n potential buyers. Buyer i pays q_i only after an accepted delivery, receives incremental benefit v_i if delivered, and incurs adoption cost a_i regardless of technical success. A shared technical core costs K to construct and qualify. Coordination and other upfront costs are A. Conditional on the core qualifying, buyer-specific delivery costs d_i are incurred; a settlement or service-route burden f_i is paid from later receipts. All quantities use the same declared numeraire and horizon.

The qualification event $Z \in \{0, 1\}$ has probability p. In this model the same event affects every buyer: after qualification, all specified adaptations succeed and pay; if qualification fails, no delivery costs or buyer receipts follow. This favorable two-state timing is an assumption for the reference model. A real proposal must add partial delivery, defaults, refunds, support failures, and their financing requirements where relevant.

For a participating set C, the expected venture contribution is

$$\Pi(C) = p \sum_{i \in C} (q_i - d_i - f_i) - (K + A). \tag{4}$$

Buyer i's expected incremental surplus is

$$U_i = p(v_i - q_i) - a_i. \tag{5}$$

These are not yet market equilibria or audited accounting statements. They are explicit participation calculations under supplied beliefs and cost classifications. The sum of buyer and venture surplus is a combined-participant measure: route-provider receipts and other outside parties are not modeled. It is a social-surplus measure only under the additional assumption that the charged construction and delivery costs represent the relevant real opportunity costs and external effects are absent. Funding charges and a required expected return can be included in the stated hurdle; they cannot be assumed away in an actual contract.

Proposition 3.1 (A shared-construction threshold)**.** *For homogeneous buyers with accepted success price q, delivery cost d, route burden f, and $p(q-d-f) > 0$, strictly positive expected venture contribution requires and is equivalent to*

$$n > \frac{K+A}{p(q-d-f)}. \tag{6}$$

The smallest integer meeting this profit condition is $n_{\min} = \lfloor (K+A)/[p(q-d-f)] \rfloor + 1$. This is not sufficient for participation, financing, rights, or technical validity.

Proof. Substitute the homogeneous quantities into Equation (4) and solve $pn(q-d-f) > K+A$. When the denominator is nonpositive and $K + A \geq 0$, no positive n yields strictly positive contribution. □

Theorem 3.1 (Joint emergence in a restricted reference arrangement)**.** *Suppose a separately developed capability incurs K for each buyer, while a qualified shared construction incurs K once. The incumbent commitment rule permits at most one binding buyer per development decision; the alternative admits n coordinated, simultaneously binding purchase commitments. If*

$$p(q-d-f) < K < pn(q-d-f), \tag{7}$$

with $A = 0$, then neither shared construction under a single committed buyer nor coordinated demand for separate constructions has positive expected contribution. Their joint change does. The result presumes the buyers' participation conditions, relevant rights, and sufficient financing.

Proof. With one bound buyer, a shared construction has contribution $p(q-d-f) - K < 0$. With separate constructions, any nonempty set of m buyers has contribution $m[p(q-d-f) - K] < 0$. With a shared construction and n commitments, contribution is $pn(q-d-f) - K > 0$. □

This is an elementary complementarity result, not a general impossibility theorem for bilateral contracting. A conventional entrepreneur allowed to coordinate commitments, forecast later sales, or invest speculatively may reach the same solution. Our unrestricted comparator is explicitly permitted to do so within the registered rules. The theorem diagnoses a restriction in the named reference formulation; it does not prove superiority over ordinary entrepreneurship.

A search that retains only individually profitable exploratory changes can miss the pair. Retaining a temporary, bounded hypothesis about a shared construction is therefore different from funding it prematurely. Exploration can allow complementary hypotheses; execution still requires a feasible, accepted package. This is closely related to the established importance of complementarities across technology and organization (Milgrom and Roberts, 1990).

3.2 Aggregate surplus does not select acceptable prices

Suppose each buyer has a maximum accepted success payment b_i. For $p > 0$, its participation constraint and payment limit imply

$$0 \leq q_i \leq \bar{q}_i = \min\{b_i, v_i - a_i/p\}. \tag{8}$$

Negative $\bar{q}_i$ excludes that buyer under this contract. The caps are elicited or assumed inputs, not measurements that a language model can declare true.

Proposition 3.2 (Conditional price feasibility)**.** *For a fixed coalition with $\bar{q}_i \geq 0$, risk-neutral participation, no other transfer constraints, and a required expected venture margin $h \geq 0$, there exist prices satisfying Equation (8) and $\Pi(C) > h$ if and only if*

$$p \sum_{i \in C} (\bar{q}_i - d_i - f_i) > K + A + h. \tag{9}$$

Proof. The venture's expected contribution is increasing in every q_i. Its maximum on the rectangular admissible price set occurs at $q_i = \bar{q}_i$. This maximum exceeds h exactly under the stated inequality. □

The result establishes feasibility under supplied information, not a fair division, truthful revelation, or equilibrium selection. Choosing every cap can leave buyers only weakly better off and fragile to error. Real negotiations can require strict margins and different risk sharing. If buyer values are private, incentive compatibility cannot be inferred from this calculation; bilateral-trade impossibility results show why no general escape from such tensions follows from more computation (Myerson and Satterthwaite, 1983).

4 A profitable plan can still be impossible to deliver

Expected profit and cash available at a date are different resources. Conditional commitments payable after acceptance do not finance earlier construction merely because the total contract value exceeds cost. The relevant constraint is the complete cash path.

Let B_0 be unrestricted authorized opening cash. Let r_t be actual receipts at stage t, c_t cash expenses, and ℓ_t a separately approved net financing flow, with repayments included when due. A plan must satisfy

$$B_t = B_0 + \sum_{s \leq t} (r_s - c_s + \ell_s) \geq B_t^{\min} \tag{10}$$

for every relevant time and required stress path. Reserved customer funds and lender commitments count only under their actual permitted uses. The inequality is a necessary monetary condition; it does not create labor, licenses, equipment, or operating capacity.

Proposition 4.1 (Prefunding under delivery-before-payment)**.** *In the two-state model, with no borrowing, zero mandatory cash buffer, all development paid first, all successful-delivery costs paid next, and route charges deducted after customer payment, the necessary and sufficient cash condition is*

$$B_0 \geq K + A + \sum_{i \in C} d_i \tag{11}$$

provided total route charges do not exceed total customer receipts. If they do, add the excess to the right-hand side.

Proof. The lowest success-path balance before customer receipts is $B_0 - (K + A) - \sum_i d_i$. The failure path spends only $K + A$. The final route deduction cannot make the success balance lower than its pre-receipt minimum unless its amount exceeds the receipts. These conditions cover every stage of the stated path. □

The financing requirement can be reduced by a genuine advance payment, authorized credit, or a supplier that accepts later compensation. Such an arrangement changes risk, rights, timing, and often cost; it is not a free consequence of the same profitability equation. This is a further potential reformulation to investigate. An AI must not silently label a conditional invoice as a deposit or treat someone else's refundable principal as its working capital.

5 Executed reference study: a shared core, four buyers, and the missing cash

5.1 What is supplied and what is computed

The laboratory is standard-library Python using exact rational arithmetic. It screens admitted subsets of supplied hypothetical buyers under supplied construction and commitment modes. A second implementation checks branch probabilities, every represented cash flow and balance, participant surplus, and the decision classification. Neither program calls a language model, invents a technical core, finds actual customers, or performs a transaction.

The main case has four buyers with $q = 300$, $v = 360$, $d = 75$, and $a = 0$; $K = 600$, $A = f = 0$, $p = 4/5$, and $B_0 = 900$. All monetary numbers in the laboratory and tables are *thousands of a common numeraire*; the prose's dollar amounts are illustrative USD interpretations, not observed revenues. Qualification is a shared event, and successful adapters meet all modeled requirements. Nonmonetary capacity is assumed sufficient. These assumptions deliberately isolate economic complementarity and cash timing.

Table 2: The four-buyer illustration. Success contribution is not expected contribution. No actual buyer or technology is represented.

Buyers	Receipts if successful	Total cost if successful	Contribution if successful	Expected contribution	Required prefunding
1	300	675	−375	−420	675
2	600	750	−150	−240	750
3	900	825	75	−60	825
4	1,200	900	300	**120**	**900**

The three-buyer plan is attractive if one reports only a successful realization. Under the supplied probability, however, its expected contribution is negative. Four buyers cross the expected-contribution threshold. They each have expected surplus $0.8(360 - 300) = 48$, for aggregate expected buyer surplus 192. Adding venture contribution 120 gives expected total surplus 312, equal to $0.8 \cdot 4(360 - 75) - 600$. Internal transfers are not counted again as new value. Interpreting 312 as economy-wide surplus requires the opportunity-cost and external-effect assumptions stated above.

5.2 The factorial comparison

The constructor evaluates all fifteen nonempty subsets of the four buyers in each of four mode combinations, retaining a stop option with zero contribution. All sixty attempted records are checked, including rejected subsets.

Table 3: Best admitted expected contribution, including the option not to build. The constrained commitment mode binds at most one buyer per development decision.

Technical construction	Single-buyer commitment	Coordinated commitments
Separate core for each buyer	0	0
Shared technical core	0	**120**

Zero means the constructor declines investment; it does not mean that performing the losing project costs nothing. A single attempted development has expected contribution −420. Four separate developments would have expected contribution −1,680 and would also exceed the main case's financing. The joint result is matched by ordinary unrestricted enumeration supplied with the same options. There is no algorithmic superiority claim.

5.3 Complete event accounting and common failure

The funded joint plan starts with 900. It pays 600 for development. If qualification fails, it stops with 300. If qualification succeeds, it pays 300 of delivery costs, then receives 1,200 following accepted deliveries, ending with 1,200. Expected terminal cash is

$$\tfrac{1}{5}\,300 + \tfrac{4}{5}\,1{,}200 = 1{,}020, \tag{12}$$

which is an expected contribution of 120 above opening cash. A successful path earns 300; the failure path loses 600. Four buyers do not turn a common 20% technical failure probability into four independent risks. An approved risk budget must be able to bear the modeled loss, and natural settings need additional downside cases.

The checker validates the two represented branches through five cash-flow equalities plus exact probability, decision, and participation checks. This certificate establishes internal consistency of the supplied model. It does not verify the truth of p, willingness to pay, technical correctness, legal enforceability, or economic independence of counterparties.

5.4 Controls that reject an attractive narrative

Table 4: Full-coalition controls. Values are in thousands; unaltered parameters equal the main case.

Change	Expected contribution	Prefunding needed	Decision or missing condition
Only 600 cash available	120	900	Reject: delivery is unfunded
One buyer uncommitted	120*	900*	Reject: assumed receipts lack commitment
Shared-use rights absent	120*	900*	Reject: sharing not admitted
Qualification probability 0.60	−60	900	Reject: expected loss
Route cost 10 per delivery	88	900	Admit within the model
Route cost 40 per delivery	−8	900	Reject: route removes viability
Upfront coordination cost 120	0	1,020	Reject: no strictly positive margin
One buyer's value below price	120*	900*	Reject: buyer participation fails

*The arithmetic is shown only to expose the failed premise. It is not a supported revenue or benefit estimate once the underlying commitment, permission, or participation fails.

The controls matter because a fluent system can quietly repair a losing spreadsheet by inventing a fourth buyer, dropping an adaptation cost, changing an output requirement, or borrowing against a nonexistent commitment. The proposed discovery process must instead identify that gap and request a valid investigation or authorized arrangement. Refusal is useful only alongside demonstrated ability to find real positive cases; rejecting everything is not the target.

5.5 Exact parameter coverage

The reference grid crosses eight buyer counts $n = 1, \ldots, 8$; five probabilities $p \in \{1/5, 2/5, 3/5, 4/5, 1\}$; seven development costs $K \in \{0, 150, 300, 450, 600, 900, 1200\}$; four per-delivery route costs $f \in \{0, 10, 40, 100\}$; and two cash settings, exactly at prefunding and one unit below it. This gives $8 \cdot 5 \cdot 7 \cdot 4 \cdot 2 = 2,240$ cases.

All constructor decisions agree with the separately implemented event-ledger checker and the closed-form profit and financing conditions. Of these cases, 540 are admitted and 1,700 rejected. These counts describe the deliberately selected parameter grid, not the prevalence of viable businesses in any economy. Relabeling or expanding such a grid would not provide independent evidence of autonomous market discovery.

The companion has 50 passing tests, including full-grid agreement, factorial modes, participation and rights flags, missing commitments, altered receipts and probability records, omitted failure paths, exact thresholds, staged reuse, and deterministic regeneration. The tests validate specified calculations and artifact semantics; no external institutional review is claimed.

6 A successor can change which later markets are viable

A retained capability may lower the discovery, qualification, or integration costs of a later offering. The relevant retained object includes its evidence, rights, and failure conditions. Reuse does not mean assuming that a previously correct method remains valid under new inputs or that previously funded costs never matter in a campaign comparison.

For a later homogeneous market, a reduction in justified setup cost from K to K' changes the profit threshold from $K/[p(q-d-f)]$ to $K'/[p(q-d-f)]$. That can make a smaller group of buyers sufficient. This is an elementary implication of Proposition 3.1, not proof that reuse will create the required reduction in a real project. Methods for learning reusable abstractions provide technical precedents, but their transfer must be measured in the proposed domain (Ellis et al., 2020).

6.1 The staged reference study

In a successful main-case exchange, cash rises from 900 to 1,200: 300 of earned contribution is available above a protected 900 reserve. Two new hypothetical buyers each offer 300 for a related capability; their delivery costs remain 75 each and the separately supplied conditional qualification probability is again 0.8.

If retained knowledge reduces the new setup cost only to 300, the second project has expected contribution $0.8 \cdot 450 - 300 = 60$, yet requires $300 + 150 = 450$ before receipts. It cannot be funded solely by the earlier 300 of earned contribution while preserving the reserve. The library has created a profitable possibility but not a financeable one under this policy.

If a separately justified reuse reduces setup to 150, required prefunding is exactly 300 and expected contribution is $0.8 \cdot 450 - 150 = 210$. The second project then fits the earned-surplus constraint. Neither cost reduction is measured technical evidence: they are alternative supplied conditions to test the accounting.

Table 5: Exact two-stage illustration with second setup 150. A second attempt occurs only after first-stage success and only with earned cash above the protected reserve.

Path	Probability	Terminal cash
First qualification fails	1/5	300
First succeeds, second qualification fails	4/25	1,050
Both qualify and deliver	16/25	1,500
Expected terminal cash	1	1,188
Expected contribution above initial 900		**288**

The three-leaf expectation is $120 + 0.8 \cdot 210 = 288$. It does not establish perpetual compounding or a market valuation. Initial risk capital was still necessary, first-stage failure can lose 600, and the second-stage probability is specified rather than learned. If the new probability is 0.2, even the reduced-cost second project is rejected. Historical success does not manufacture confidence in a different construction.

6.2 The commercially meaningful archive

The archive should retain a common core only with a declared domain, output obligations, licenses, environment dependencies, adaptation procedure, independent checking path, complete lifecycle costs, and

invalidation conditions. A project history records what failed as well as what worked. A router is responsible for deciding whether a new problem falls inside the supported domain; component correctness does not prove router correctness.

The strategic objective is not maximizing the number of stored services. It is expanding the set of economically and technically defensible outcomes that can be delivered. Some useful successor changes delete a fragile dependency, retire a misleading abstraction, or reveal that two apparently similar markets cannot safely share a capability. That negative knowledge can improve future selection without producing immediate revenue.

7 How to search for an economy that does not yet exist

7.1 Begin with a restriction map, not a business-plan template

A candidate opportunity starts with evidence of a required outcome and a named reason it is currently unserved. The restriction might be a costly inference, a repeated proof, a nontransferable representation, incompatible delivery conditions, an unfunded interval, or buyer fragmentation. Some restrictions are essential; others are inherited means. Classification must be open to challenge without giving the candidate authority to erase obligations.

Represent the opportunity as a typed dependency structure. Outcome nodes connect to required capabilities; capabilities connect to construction experiments, resource profiles, and checkers; accepted purchase terms connect to financing and delivery obligations. A hyperedge can represent a combination whose benefit exists only when all prerequisites are present. Unlike ordinary allocation of a fixed catalogue, proposed new capabilities can introduce new nodes and realizations. Combinatorial winner determination is a useful subroutine, not the entire discovery problem (Sandholm, 2002).

An infeasible candidate should produce a specific obstruction: insufficient committed contribution, a rights conflict, a missing correctness bridge, or an unfinanceable cash prefix. That obstruction becomes a question for further research. “Find three more buyers” is not a solution unless they exist and accept; “reduce setup cost” is not a construction unless a method can do so. The system must earn those changes through evidence.

7.2 Search jointly over capability and arrangement

The proposal space includes a new representation or general construction, a discriminating experiment, a beneficiary set, delivery adaptations, contract timing, a risk-bearing arrangement, and a supported operating route. A candidate can also use an existing solution, stop, or request clarification. No objective should reward the number of reformulations irrespective of their outcomes.

The technical generator can use model proposals, symbolic synthesis, automated program search, and installed solvers. The economic generator can suggest bounded combinations of counterparties and reviewed contract templates. A feasibility checker can reject proposals before expensive construction, but a supplied cost estimate remains a hypothesis. The natural-language model cannot be both the source of every uncertain input and the authority certifying their truth.

Searching fixed templates is a first engineering step, not a complete transformational system. New technical operators or commercial forms may eventually be proposed and admitted through separate review. Their construction, discovery cost, new risks, and qualification must remain visible. Otherwise the fixed operator catalogue becomes the next unnecessarily restrictive formulation.

7.3 Create decisive evidence before making irreversible commitments

An early experiment should distinguish mechanisms, not merely produce an attractive score. A proof obligation might determine whether several kernels share a common invariant. A small adaptation test can reveal whether apparent reuse hides per-customer redevelopment. An authorized discussion can determine

whether a buyer needs the proposed artifact, a stronger audit trail, or a different service window. A financing rehearsal can expose an unfunded prefix before anyone spends money.

Advance market commitments and related arrangements show that demand can be organized around a not-yet-available capability (Kremer et al., 2020). They do not make buyer acceptance automatic or eliminate development finance. The system should treat confirmed terms, nonbinding interest, estimated value, and executed payment as distinct evidence states.

Algorithm 1 Joint capability–arrangement discovery (proposed)

Require: Authorized purposes, public task evidence, retained capabilities, bounded discovery resources
Reproduce the incumbent outcomes, costs, contracts, and cash timing
Identify a restriction and distinguish requirements from inherited means
Propose missing capabilities and combinations of independently motivated beneficiaries
Construct candidate methods and mechanism-discriminating tests
Propose permitted arrangements; keep unaccepted terms explicitly conditional
for each selected exploratory package within budget **do**
    Check rights, required outcomes, participant constraints, resources, and cash prefixes
    Execute authorized technical tests; obtain actual counterparty evidence separately
    Retain costs, failures, unknowns, and complementary unfinished alternatives
end for
Freeze the finalist, incumbent, unrestricted comparator, and evidence protocol
Obtain separate approvals for funding, agreements, and any operational execution
Perform the admitted construction, independent acceptance, delivery, and settlement
Compare realized outcomes and complete economics; keep price speculation separate
Test a fresh independent adoption and revalidate any retained capability
Promote discovery-policy changes only after separate complete-campaign evaluation

7.4 No automatic victory from complementary edits

A jointly changed package may require more development time, more counterparties, and more legal or integration work. Those are costs of the transformation, not exceptions to its budget. A strong unrestricted baseline must be allowed to discover the same arrangement. A useful result may be that the system independently reconstructs a known valuable method; that is a discovery achievement relative to its starting information, not scientific novelty.

The prototype in Section 5 supplies the capabilities and economic conditions. It therefore tests decision logic, not the difficult invention steps in Algorithm 1. The decisive empirical advance is an independently selected natural workload where the system constructs a previously unspecified method and obtains real, independently acceptable terms that outperform credible alternatives after full costs.

8 The abundance paradox: do not monetize the work intelligence should eliminate

A token mechanism that charges only for reasoning steps can conflict with a system that learns to avoid those steps. The same applies to a service provider selling labor hours when a general construction removes most of the labor. This is not an argument against metered services; it is a reason to choose the economic unit deliberately.

Proposition 8.1 (A conditional abundance paradox)**.** *Suppose a technical productivity factor $g \geq 1$ reduces required operations per accepted outcome from c to c/g. Suppose accepted demand becomes $N_0 g^{\varepsilon}$ and the fee per operation remains τ. Then aggregate operation-fee flow is*

$$F_{\mathrm{op}}(g) = \tau c N_0 g^{\varepsilon - 1}. \tag{13}$$

It decreases with productivity when $\varepsilon < 1$. A fixed fee φ per accepted outcome instead gives $F_{\mathrm{out}}(g) = \varphi N_0 g^{\varepsilon}$ under the same demand assumption.

Proof. Multiply the stipulated outcome count by operations per outcome and the fee per operation. The sign follows from the exponent. The outcome-fee expression follows by direct multiplication. □

The result is conditional on fixed fees and the specified demand response. Competition, better alternatives, changing quality, and falling willingness to pay can alter either fee. It is not a universal argument that value-based fees grow forever. With illustrative $\varepsilon = 1/2$, a sixteenfold productivity improvement multiplies accepted outcomes by four but divides operation-fee flow by four. A fixed outcome fee avoids that *mechanical* decline, not every market risk.

The proposed commercial unit is therefore an independently valuable capability or bounded service commitment: a qualified implementation, an accepted investigation, a supported solver, or a maintenance period. Compensation is for what is actually delivered under competitive terms. A buyer should not be charged for imaginary counterfactual labor, and an openly licensed algorithm does not become subject to perpetual royalties because a registry mentions it.

This aligns the technical search with the service relationship. A deeper transformation can leave the buyer better off while making delivery less costly. Whether providers capture some of that benefit depends on actual competition and accepted prices. The original market is allowed to shrink when its work becomes unnecessary; the system then searches for newly feasible outcomes rather than fabricating additional tasks.

9 $AGIALPHA: creating value and capturing demand are different problems

9.1 Three ledgers that must remain separate

The project-specific motivation is a more useful native economy, not a promised market price. The retained design describes services, commitments, and settlement while explicitly withholding parent-company equity, automatic dividends, or infrastructure ownership (MONTREAL.AI, 2026b). The pinned repository README also distinguishes AGIJobManager's USDC route from Alpha / Jobs v0's $AGIALPHA routes (MONTREAL.AI, 2026a). This paper changes none of those terms.

Maintain three separate ledgers: customer outcomes, venture receipts and obligations, and native token operations. A transfer can be a venture receipt without being net profit; a profitable venture can provide no token-holder cash-flow right; a useful native payment can be immediately sold by its recipient. Transaction volume, standing token demand, and market capitalization therefore need separate models.

Token-valuation research already analyzes adoption and transactional demand rather than assigning all corporate profit to a utility token (Cong et al., 2021). Here the research question is narrower and more directly testable: does a constructed service create independently chosen native activity whose coordination benefits exceed its burdens? A high valuation remains a possible market outcome, not a conclusion of the model.

9.2 A route can destroy the market it is meant to support

The economic mechanism cannot be treated as a free layer added after a useful business has been found. Its settlement, conversion, review, reserve, and operating costs can eliminate the very surplus that makes participation possible.

Corollary 9.1 (A strict ceiling on an additional route burden)**.** *In the homogeneous shared-construction model, with the success price and other conditions fixed, positive expected contribution requires*

$$f < q - d - \frac{K + A}{pn}, \qquad p > 0. \tag{14}$$

If the right-hand side is nonpositive, no nonnegative additional route burden can make that package strictly profitable without another change.

Proof. Rearrange Equation (4). This condition is only the venture-margin constraint; buyer, funding, and rights constraints can be tighter. □

In the main synthetic case, the strict ceiling is $300 - 75 - 600/(0.8 \cdot 4) = 37.5$ per successful delivery, in thousands. An extra burden of 10 leaves expected contribution 88. A burden of 40 changes it to −8. These amounts are not actual protocol fees or recommended prices. The variable represents an explicitly added total cost. Its purpose is to show that extracting more from a newly viable market can make it cease to exist.

A better native route could also reduce a genuine coordination or enforcement cost. That potential benefit must be measured against a matched alternative, not presumed because the service uses a blockchain. Platform economics emphasizes that participation on the different sides and the allocation of charges matter jointly (Rochet and Tirole, 2003). The product should be easier or better to deliver, not burdened merely to produce token throughput.

9.3 Settlement demand is not a valuation formula

For a stationary illustrative flow of accepted services with annual numeraire value X and average holding time h days, a transaction-inventory approximation is

$$B_{\text{tx}} \approx Xh/365. \tag{15}$$

It is an accounting heuristic for average transaction balances in the same value units, not a market-capitalization estimate. Separately committed operating balances may be added only when they do not overlap with the inventory already counted. The duration, concentration, liquidity, and asset exposure of actual obligations must be observed.

At hypothetical $X = \$1$ billion, one day of holding corresponds to about \$2.74 million of average transaction inventory; thirty days corresponds to about \$82.19 million. Neither number supplies a price floor. Artificial delays introduced solely to enlarge this calculation would make the service worse and need not improve equilibrium demand.

Proposition 9.1 (Service success alone does not identify a token price)**.** *Suppose service prices are set in an external numeraire, customers can acquire the required token immediately before settlement, recipients can sell immediately after settlement, and holders have no claim to service-provider profits. Service quantities, numeraire revenue, and provider profit alone do not identify the token's unit price.*

Proof. For any stipulated positive exchange price P, a service priced at q uses q/P token units. Changing P and rescaling the units transferred preserves the service quantities, numeraire payments, and provider profit. These operating observations therefore do not select a unique P without additional supply, portfolio-demand, liquidity, institutional, or equilibrium restrictions. □

This is an identification statement, not a claim that all such prices are market-clearing equilibria or that token utility is irrelevant. It explains why the observed success of an AI business does not itself establish an extraordinary token valuation. Restrictions on minting, administrative powers, concentration, unlocks, market depth, and holder rights would matter in a full analysis; no live token audit or price estimate is performed here.

9.4 A defensible native hypothesis

The proposed native roles are concrete services: commissioning admitted work, settling accepted outcomes, and supporting appropriately scoped provider commitments or capacity-backed reservations where separately implemented and authorized. Each role must name what a participant receives, who is responsible, the asset exposure, the duration, and the conditions for release, cancellation, or recovery. Existing protocol behavior must not be silently reinterpreted to create those products.

Burning tokens is a supply operation, not a source of venture cash; destroyed units cannot also finance delivery. A positive market-price response to a burn is not implied by arithmetic alone. Idle holdings, circular settlements, and internal demonstrations are not substitutes for external adoption. The empirical test is whether users and providers choose useful native services on defensible terms, not whether promotional activity temporarily changes a quoted price.

A mature evidence case would show a growing set of independently deliverable outcomes, repeat use, sustainable provider economics, transparent native balances tied to actual obligations, and a competitive reason for using the route. It would not promise that every new capability transfers wealth to every token holder. Calling an instrument a utility token is also not a legal determination; Canadian guidance expressly discusses securities-law implications for such offerings (Canadian Securities Administrators, 2018). New financial rights, fundraising arrangements, or promotional representations require separate qualified review.

10 SUCCESSOR Ω: the implementation target

10.1 Extend the existing mission, not a parallel speculative product

The named implementation target is

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

Find your Alpha. Build your SUCCESSOR Ω.

The pinned README at commit `b44ba296e00500c91c1617fc942b5f0050ceabd6` describes opportunity investigation, strategy development, problem reframing, configured Computer Use, and retained project history. It distinguishes preparation, native execution, review, and settlement (MONTREAL.AI, 2026a). We cite that snapshot as an implementation starting point, not as an assertion about all subsequent releases or evidence of automated commercial formation.

Find your Alpha becomes the identification of an economically meaningful removable restriction. *Build your SUCCESSOR* Ω becomes the construction and retention of a capability that changes what subsequent missions can deliver. Alpha supplies discovery and development; Jobs and Fleet retain their respective commissioned-work and qualified-execution boundaries. A connected interface does not itself create a customer, a valid contract, or permission to spend.

10.2 One proposed opportunity record

The next additive schema should keep one mission identity while separating technical proposals, modeled economics, actual counterparty evidence, operational approval, delivery, and reuse. Proposed components are implementation boundaries, not files claimed to exist:

The reference certificate uses the label `synthetic-model-only`; its `live_acceptance` and `authority_granted` fields are false. A reader can verify model arithmetic without turning the file into an authorization token. Production integration needs its own schema review, source qualification, actual domain checkers, and participant commissioning. No merge, release, deployed contract, token supply, fee, holder right, ENS configuration, or wallet action is changed by this manuscript.

Table 6: A proposed versioned extension to the existing runtime.

Component	Responsibility
`economic_possibility`	Bind purposes, technical candidates, beneficiary requirements, uncertainties, and proposed arrangements
`capability_commonality`	Investigate shared structure and validate every required adaptation
`arrangement_feasibility`	Check supplied participation, funding, rights, resource, and timing constraints
`counterparty_evidence`	Distinguish interest, proposals, accepted terms, and actual performance; imports grant no authority
`realization_ledger`	Reconcile costs, liabilities, receipts, and accepted outcomes in consistent units
`native_route_comparison`	Measure the native mechanism's added benefit and burden without changing existing assets
`successor_market_reuse`	Retain conditional capability and delivery evidence; test fresh independent adoption

10.3 A user experience centered on outcomes

The interface should begin with a beneficiary's intended outcome and approved public task material, not a blank form asking the user to invent the winning strategy. It should present the current restriction, proposed common capability, what each participant must receive, what is known, what remains uncertain, and the next bounded test.

After a technical experiment, the user sees actual results rather than a self-assessed success message. After a commercial step, the record distinguishes a received offer from a signed commitment and a commitment from collected money. English and French views should expose the same terms and uncertainty. Public case studies should disclose only information authorized for publication; the native public-input scope does not authorize access to confidential customer systems or personal records.

The product should also make it easy to decline. A proposed fourth customer without evidence, an unfunded delivery interval, an incompatible license, or a model that loses to a standard solver should stop the associated claim. This discipline protects the ability to make strong claims when the complete package does work.

11 An empirical program that can reject the thesis

11.1 The unit is the complete opportunity campaign

The experimental unit includes initial evidence, technical investigation, commercial elicitation, candidate construction, finalist selection, fresh evaluation, accepted terms, delivery, and the proposed reuse horizon. Several customers depending on one technical core do not become independent demonstrations of discovery. Replication should occur across independently initialized campaigns, with task families and actual counterparties selected without planting the desired transformation.

For a registered campaign initialization Z, discovery policy η, and total budget B, let $W(\eta, Z, B)$ report a vector of accepted outcomes, customer surplus or defensible benefit, realized venture contribution, complete costs, violations, and native-route observations. A scalar objective may be used only after its units and trade-offs are authorized. Model-score improvements and market-price movements should not silently replace these outcomes.

Adaptive selection creates a risk that repeatedly consulted evidence becomes a development set rather than a confirmation set. Freeze finalists and comparisons before fresh evaluation, retain unsuccessful attempts,

and use statistical methods appropriate to the actual adaptive design (Dwork et al., 2014; Howard et al., 2021). No confidence interval repairs fabricated demand or an invalid outcome instrument.

11.2 Comparators that make joint invention a meaningful claim

Cross capability mutability with arrangement mutability. A technology-only arm may optimize the method while retaining the registered commercial arrangement; an arrangement-only arm may optimize terms around an incumbent capability; the joint arm can change both. The restricted arms diagnose mechanisms rather than serve as the only competition.

Include a strong unrestricted agent explicitly allowed to coordinate, reframe, reuse, and propose new arrangements, as well as competent human technical and commercial teams under a declared resource accounting. Give all arms the same permissible tools and initial evidence. Automated agent/workflow search such as ADAS, AFlow, and GEPA supplies additional relevant comparator families (Hu et al., 2024; Zhang et al., 2024; Agrawal et al., 2025). Budget quality matters: a larger unseen engineering effort must not masquerade as a better search architecture.

Supplying the discovered method or arrangement to another arm is a useful diagnostic. It separates ability to discover the alternative from ability to execute it. If every method finds the same package, the market can still be valuable, but no superiority of the proposed discovery policy follows. If the package only works because a counterparty supplied extra resources after seeing the expected winner, that difference must be included.

11.3 Six falsifiable hypotheses

H1: joint formation. At matched total resources, joint discovery produces more independently accepted, technically valid, economically realizable opportunities than strong unrestricted alternatives. The claim fails if the gain disappears after full discovery, integration, legal, review, and operator costs.

H2: structural attribution. A named technical and commercial complementarity accounts for a material part of the gain. Appropriate ablations or diagnostic controls weaken the result when either change is removed. The claim fails if an unchanged-frame method achieves the same outcome on comparable terms.

H3: funding realism. The system distinguishes profitable forecasts from feasible delivery cash paths and correctly handles delayed receipts, partial delivery, failure, and refund obligations. Counting conditional orders as unrestricted cash falsifies the claim.

H4: transferable commercial capability. A retained capability supports fresh independent adoption under preserved requirements and complete lifecycle costs. Familiar labels, copied customers, or a superior inherited starting artifact alone do not demonstrate better discovery.

H5: abundance-compatible economics. New constructions reduce unnecessary work without making the service economy depend on preserving that work. The claim fails when apparent provider gains come from charging for undelivered activity or ignoring competition and customer costs.

H6: native-route contribution. On matched services, the admitted $AGIALPHA route provides measurable coordination or service value sufficient to justify its total burden. The claim fails when equivalent nonnative delivery is consistently better after all costs, or when volume is mainly internal or subsidized activity. This tests a mechanism, not a token-price target.

11.4 A decisive first natural study

Select several independently documented, public-input computational needs with overlapping structure but distinct required outputs. Each beneficiary states an affordability limit, acceptance criteria, delivery constraints, and a lawful route for a bounded offer. A research campaign investigates whether a shared method can be constructed without being handed that method in advance.

Before substantial execution, agree on who bears failure, how costs preceding delivery are financed, the exact scope of sharing, and what happens if one buyer leaves. Produce and independently test the capability. Execute one admitted delivery and payment. Then test an independent adoption of the retained method. Publish actual results and limits, including a failed or declined opportunity selected by the same process. A negative result can be a valid experiment, but not evidence that the market has been created.

12 Practical design choices and unresolved risks

12.1 Counterparties are agents, not entries in a spreadsheet

Willingness to pay is contextual and may change after a prototype exists. Buyers can withdraw, understate value, wait for others to fund an open-source result, or require exclusivity incompatible with sharing. Suppliers can promise capacity they cannot deliver. These are not implementation details solved by model confidence. Accepted terms, rights, funding, and operational capacity need independent evidence at the relevant time.

Our exact checker assumes complete, truthful model inputs. It is neither an incentive-compatible procurement mechanism nor a substitute for contracting and competition-law review. Grouping independent requirements to fund a shared technical result should not become coordinated manipulation of downstream prices or restriction of customer choice. A proposed arrangement must pass the applicable professional and legal review before use.

12.2 Common capabilities create common exposures

A shared core can reduce fixed cost while correlating failures and concentrating operational dependence. The reference's single qualification event makes that correlation explicit; natural service failures may be more complex. Resilience, fallback providers, independent checking, migration, and retirement are part of complete delivery cost. A common registry must not manufacture trust by counting several signatures controlled by one party as independent review.

Portfolio claims must distinguish separate technical risks from renamed copies of the same risk. Nor should reuse become a way to amortize a past expense twice. Campaign comparisons need matched initial libraries or explicit inherited-asset valuation, while realized ledgers count each external receipt once and consolidate internal transfers.

12.3 The evaluator can be restrictive, but cannot be privately replaced

An evaluator that demands the original workflow can penalize a better permitted result. An evaluation instrument may therefore need revision. That revision requires evidence that it better measures the authorized outcome, independently adjudicated disagreement cases, and a common comparison of incumbent and challenger. A candidate cannot select an easier judge and declare its own commercial success.

Technical correctness, customer usefulness, commercial contribution, and native demand require different instruments. A proof of numerical equivalence does not prove customer benefit; a buyer's acceptance does not establish a new scientific theorem; a token receipt does not establish profit. The architecture depends on keeping those distinctions operational rather than burying them in one confidence score.

12.4 Ambition without a fabricated endpoint

The long-horizon ambition is a network that repeatedly helps people realize valuable capabilities that were previously unavailable on acceptable terms. A later capability can support more demanding science, production, or infrastructure. Progress along that path requires actual commissioned resources and legitimate cooperation, as the preceding production-frontier paper emphasizes (Boucher, 2026c).

No finite reference study establishes exponential intelligence, indefinite profitability, civilizational control, or extreme asset appreciation. Those conclusions would require assumptions and evidence absent here. The present objective is still substantial: make the creation of viable new services a learnable, executable, independently testable activity rather than only a human-authored business narrative.

13 Relationship to prior work

Transforming conceptual spaces. Boden's distinction between exploring and transforming conceptual spaces, and Ritchie's demand for testable definitions, directly motivate the formulation-level question (Boden, 1998; Ritchie, 2006). Our economic-possibility package names the technical and institutional objects that change and the evidence that could refute their alleged benefit.

Innovation and complementary organization. Endogenous-growth theory treats knowledge creation as an economic activity (Romer, 1990); complementarity models connect technology, strategy, and organization (Milgrom and Roberts, 1990). Our finite results are not replacements for those theories. They supply transparent local conditions and operational evidence categories for an agent-based research program.

Markets, commitments, and strategic information. Combinatorial allocation, advance market commitments, and two-sided platform models provide mechanisms for coordinated activity (Sandholm, 2002; Kremer et al., 2020; Rochet and Tirole, 2003). We add capability construction to the search object rather than assume the catalogue is complete. Myerson–Satterthwaite remains a reminder that more intelligent search does not remove the economic limits imposed by private information and incentive requirements (Myerson and Satterthwaite, 1983).

Executable invention and automated design. DreamCoder learns reusable abstractions, AlphaEvolve searches executable algorithms, and ADAS, AFlow, and GEPA search aspects of agent or workflow design (Ellis et al., 2020; Novikov et al., 2025; Hu et al., 2024; Zhang et al., 2024; Agrawal et al., 2025). These are technical precedents and competitive baselines. Depending on their actual behavior, they could implement parts of the proposed joint discovery; the category is not determined by a project name.

Evidence and token economics. Adaptive-data and sequential-inference work guide confirmation design (Dwork et al., 2014; Howard et al., 2021). Tokenomics supplies a framework for understanding transactional adoption and valuation (Cong et al., 2021). Our route-cost and abundance calculations are conditional accounting results, not a fitted token-pricing model. The project-specific economic boundaries come from the retained design and pinned repository, not a new entitlement or legal opinion (MONTREAL.AI, 2026b,a; Canadian Securities Administrators, 2018).

14 Conclusion: invent what makes a market possible

The next valuable output of intelligence need not be another answer inside an existing market. It can be a shared construction and a voluntary arrangement that make previously unserved needs worth serving. That requires intelligence applied to methods, experiments, customer boundaries, financing, delivery, and the unit of compensation together.

The exact study makes the distinction concrete. Neither fragmented demand with a shared core nor coordinated demand for separate cores supports a viable investment under the stated rules. Together they do. Yet positive expected contribution still fails without delivery finance, and a sufficiently expensive native route destroys the opportunity it was intended to serve. Reuse can lower a later market's threshold, but its next cash path must also be funded. These results are simple, auditable, and conditional; the hard research challenge is to discover the corresponding structure in natural tasks and real voluntary arrangements.

For SUCCESSOR Ω, the ambition is to retain capabilities that improve what later missions can achieve. For $AGIALPHA, the economically relevant test is whether valuable native service relationships emerge and remain competitive, not whether business profit can be relabeled as holder wealth. Keeping this boundary explicit strengthens the research question rather than diminishing it.

Do not only build intelligence that competes inside the existing economy. Build intelligence that discovers the capability and the arrangement that make a new economic possibility realizable—and earns the evidence that it works.

Research status and reproducibility

This manuscript reports an original synthesis, elementary conditional results, and newly executed exact synthetic studies. The reference constructor screens supplied alternatives; it does not discover a real technical core or establish customer demand. No frontier-model campaign, real commercial agreement, payment, live deployment, or token-price experiment is claimed. The separate checker shares authorship and is not an independent institutional audit. The companion contains source, exact records, 50 passing tests, proposed pilot instructions, and a claim–evidence register. The earlier papers and repository are unchanged.

A Reproduction and interpretation of the reference artifacts

The companion uses Python's standard library; the recorded runtime is Python 3.13.5. No credentials, network, model provider, wallet, or external account is required.

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

The driver regenerates five deterministic JSON files. Amounts use thousands of a common numeraire; rationals serialize as numerator–denominator pairs. The 2,240-case grid contains its complete input combinations and outcomes. The 60 factorial records are checked during generation; the published summary reports all four optima and counts. The separate checker recomputes the main record from event flows rather than calling the constructor's decision function.

Table 7: Included artifacts and their scope.

Artifact	Meaning
`experiments/market_model.py`	Supplied-model subset screening, thresholds, and staged reuse
`experiments/check_certificate.py`	Separate cash-path, probability, participation, and decision checker
`experiments/run_study.py`	Deterministic results driver
`tests/test_study.py`	50 tests, including the complete parameter grid and artifact mutation checks
`results/formation_certificate.json`	One exact model record; no real acceptance or authority
`results/formation_results.json`	Coalition table, factorial, risk record, and negative controls
`results/exhaustive_grid.json`	All 2,240 reference cases
`results/reuse_results.json`	Profitable-but-unfunded and funded second-stage paths
`results/token_accounting.json`	Illustrative inventory and abundance arithmetic; no price forecast
`PILOT_PROTOCOL.md`	Proposed natural-study handoff and stage-specific evidence
`CLAIMS.md`	Claims, their support, and their limits

The LaTeX file embeds its tables and bibliography and requires no external figure files or BibTeX pass. Compile with `latexmk -pdf Invent_the_Economy.tex` or run `pdflatex` twice. A file digest binds bytes, not factual truth, lawful consent, market independence, or the completeness of the modeled obligations.

B A minimum natural-campaign contract

The pilot should identify actual prospective counterparties before making demand claims. Its task inputs remain public, synthetic, or explicitly licensed and nonpersonal within the admitted workflow. Any additional confidential commercial handling needs separate approval and an appropriate system; it is not implied by this paper.

The first implementation milestone is one complete admitted campaign, not a claim that the campaign must succeed. The strongest positive outcome would show that a previously unspecified construction and accepted arrangement jointly enabled useful work, followed by independent reuse. The strongest negative outcome would identify a genuine obstruction without inventing a counterparty, lowering a required standard, or hiding a cost to preserve the story.

Table 8: Required fields before the corresponding stage may advance.

Boundary	Minimum record
Purpose	Each beneficiary's required outcome, quality, timing, rights, and unacceptable side effects
Baseline	Best credible technical and commercial alternatives; observed costs or explicit unknowns
Construction	Candidate core, adaptation obligations, proposed experiments, full discovery budget
Demand	Source of interest, nonbinding versus accepted status, actual authorized terms and expiry
Risk	Failure states, responsible parties, maximum exposure, refund and recovery obligations
Finance	Every pre-receipt cash expense, reserves, permitted funding source, and staged commitments
Delivery	Qualified environment, independent checker, acceptance process, evidence and actual receipts
Native route	Exact asset, supported contract path, fees, conversion exposure, and no inherited authority
Reuse	License, compatible conditions, fresh beneficiary, cost and correctness checks, retirement triggers
Publication	Consent, redactions, claim scope, exact source/artifact identities, and negative results

C Claim–evidence register

Claim	Support	Boundary
Technical and commercial changes can be complementary	Theorem 3.1; exact factorial	Supplied construction and commitment rules, not universal business superiority
Four buyers yield expected contribution 120 and successful contribution 300	Exact rational calculations and event-ledger checker	Assumed prices, costs, probability, and all-success adaptation semantics
Profitable does not imply financeable	Proposition 4.1; 600-cash rejection	No borrowing; explicit delivery-before-payment timing
Additional route burden can eliminate viability	Corollary 9.1; 10 and 40 controls	Fixed success prices and other model inputs; not actual protocol pricing
Reference decisions agree across the grid	2,240 cases; 540 admitted, 1,700 rejected	Finite synthetic grid, not a market prevalence estimate
Reuse can enable a funded successor offering	Three-leaf staged model; expected contribution 288	Cost reductions and conditional success probabilities are supplied
Higher productivity can reduce operation-fee flow	Proposition 8.1	Fixed fees and stipulated demand elasticity
Service success alone does not determine token price	Proposition 9.1	Identification argument, not an equilibrium or valuation model
SUCCESSOR Ω can be the implementation target	Pinned project README and proposed schema	No runtime integration, release, or commissioning performed
AI can discover and realize new natural markets	Hypotheses H1–H6; proposed protocol	Unestablished; requires real construction, counterparties, delivery, and replication

References

Lakshya A. Agrawal et al. *GEPA: Reflective Prompt Evolution Can Outperform Reinforcement Learning*. arXiv:2507.19457, 2025. https://arxiv.org/abs/2507.19457.

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, 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, 2026. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/main/docs/Beyond_the_Original_Problem.pdf.

Vincent Boucher. *Beyond the Production Frontier: Transformational Intelligence from Verified Discovery to Stellar-Scale Civilization*. MONTREAL.AI research manuscript, 2026. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/main/docs/Beyond_the_Production_Frontier.pdf.

Vincent Boucher. *Earn the Next Experiment: Transformational Intelligence through an Improvement-for-Capacity Pilot*. MONTREAL.AI research manuscript, 2026. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/main/docs/Earn_the_Next_Experiment.pdf.

Canadian Securities Administrators. *Canadian securities regulators provide additional guidance on securities law implications for offerings of tokens*. June 11, 2018; guidance introducing CSA Staff Notice 46-308. Accessed October 2026. https://www.securities-administrators.ca/news/canadian-securities-regulators-provide-additional-guidance-on-securities-law-implications-for-offerings-of-tokens/.

Lin William Cong, Ye Li, and Neng Wang. Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 34(3):1105–1155, 2021. https://doi.org/10.1093/rfs/hhaa089. Open working-paper record: https://www.nber.org/papers/w27222.

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.

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.

Michael Kremer, Jonathan Levin, and Christopher M. Snyder. Advance market commitments: Insights from theory and experience. *AEA Papers and Proceedings*, 110:269–273, 2020. https://doi.org/10.1257/pandp.20201017.

Paul Milgrom and John Roberts. The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review*, 80(3):511–528, 1990. https://www.jstor.org/stable/2006681.

MONTREAL.AI. *Find your Alpha. Build your SUCCESSOR Ω.* AGI-Alpha-Agent-v0 README at commit `b44ba296e00500c91c1617fc942b5f0050ceabd6`; source blob `1f65ee5a12cc36891dae9bab76424ff027b8105b`. Pinned source inspected for this manuscript. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/b44ba296e00500c91c1617fc942b5f0050ceabd6/README.md.

MONTREAL.AI project working material. *$AGIALPHA × SUCCESSOR Ω: A Productive-Capability Economy, Not a Price-Dependent Funding Loop*. October 19, 2026. Retained source: `AGIALPHA_SUCCESSOR_Capability_Economy.md`, supplied with the preceding project discussion and included unmodified in the companion's `sources/` directory. A proposed design, not a deployed rights instrument or legal opinion.

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

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

Graeme Ritchie. The transformational creativity hypothesis. *New Generation Computing*, 24:241–266, 2006. https://doi.org/10.1007/BF03037334.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003. https://doi.org/10.1162/154247603322493212.

Paul M. Romer. Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102, 1990. https://doi.org/10.1086/261725.

Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1–2):1–54, 2002. Open author version: https://www.cs.cmu.edu/~sandholm/oralg.aij.pdf.

Jiayi Zhang et al. *AFlow: Automating Agentic Workflow Generation*. arXiv:2410.10762, 2024. https://arxiv.org/abs/2410.10762.