USEFUL WORK ⟶ SURPLUS ⟶ GREATER CAPABILITY

AGI ALPHA FLYWHEEL

The Economic Engine

Useful Work, Collected Surplus and
Bounded Capability Investment

SUCCESSOR Ω HARNESS

Powered by TRANSFORMATIONAL INTELLIGENCE ∞

A separate financial and source-mechanism amendment
preserving the completed human-work economic comparison

Prepared with AI assistance for scientific and repository review.

Abstract

An enterprise can fund further useful construction when accepted work produces collected revenue above complete costs and required reserves. Proof of an accelerating improvement process is not necessary for that ordinary economic mechanism. Its existence for a particular product, however, depends on price, reliability, demand, attention, fixed costs, liquidity and the usefulness of the next investment. We extend a preserved native study and market-data-grounded operating-cost comparison with a separately preregistered cash model: 144 primary paths, 32 demand sensitivities and 27 source-specific settlement scenarios. The primary evidence-pack profile retains the prior complete-cost assumptions; a separate USD0.50 delivery/ USD200 human-replacement profile is an explicit unvalidated hypothesis. Four allocation ceilings range from zero to 75% of recovered, collected after-tax operating surplus. Construction failures, reserves, uncollected invoices and unsuccessful or uneconomic investments remain charged or restricted. At USD20 per accepted output, the evidence-pack base incurs a modeled USD 2,630.00 provider loss over the operating horizon. The separate fifty-cent hypothesis generates USD 121,120.44 after the tax reserve before discretionary investment, while its complete provider expense is USD 2.61 per expected accepted output. Its 50% policy earmarks USD 60,560.22 but spends zero because the frozen investment gate rejects the remaining-horizon payback. At fourfold assumed demand, a stipulated successful candidate can produce a positive simulated return after one failed construction; the all-failed world loses the construction budget. Source inspection separates Manager USDC from applicable AGIALPHA routes and identifies the actual recipient, burn, collateral and authority boundaries. No simulation is reported as a sale, a live transaction, a discovered capability or proof of self-acceleration. Authenticated executor and live financial execution remain **NOT RUN**.

Measured native components Published labor prices Explicit financial assumptions
Customer savings, company cash and token-holder outcomes have separate accounts.

1 The customer buys useful work

A project coordinator needs a readable source brief and a short presentation before publishing a paper and planning executor integration. The accepted native example supplies both files, source identities, exact requested passages and machine-readable provenance. Its retained module can be imported into a changed mission. A competent person can produce the same bounded deliverable with ordinary non-AI tools, templates, search and professional judgment. The commercial decision concerns the accepted output, total cost, waiting time and remaining attention.[1–3]

The measured example remains modest: two approved documents and four supplied passages. Its local execution and acceptance boundary took 0.111703 seconds, with zero task-time model calls. This is a native component measurement. Whole service cost, field supervision, customer acceptance and commercial proceeds were not measured. The prior paper therefore added dated labor/model/infrastructure prices and explicit operational assumptions, including failed attempts, setup, integration, reviews, maintenance and capability construction.[2]

In its primary U.S. office scenario, each route's own USD25,000 aggregate resource ceiling gave USD31.74 per expected accepted Agent output, USD28.52 for traditional completion at an assumed thirty-minute productive duration, and USD35.54 for complete-demand human fallback. These values do not establish a commercial win for that brief. Longer assumed traditional durations can change the comparison; no matched human duration was observed. The present financial analysis preserves that evidence and adds the supplier's price and cash boundary.[2]

The 42-job study's specialist results answer a different question: which internal method should be retained? Its planning repair lost all twelve fresh comparisons against the available specialist, and incremental source repair did not show a CPU advantage. Those results support retaining the stronger method. They do not determine the price or quality of traditional human work. Similarly, a positive financial simulation cannot reverse a failed method comparison.[1]

2 An explicit economic engine

For an accepted job priced at USD20, a hypothetical complete variable delivery cost of USD0.50 leaves USD19.50 *before* further business costs and tax. Half of that provisional amount is USD9.75. A customer's hypothetical USD180 saving relative to a USD200 human alternative is not company revenue. The executable model pays fixed enablement, mandatory improvement upkeep, acquisition, billing, administration, refunds and other stated costs, then restricts reserves before any new appropriation. The USD0.50 profile is a hypothesis with its own assumed demand and attention; it is not the observed cost of the four-passage brief.

Let N_m denote requested jobs served in month m, q the assumed acceptance probability, P the accepted-delivery face price, and g the collected fraction. Collection occurs one month after delivery. With refund fraction r, percentage processing charge a and fixed collection charge f, net collected cash is

$$X_m = gqN_{m-1}\{P(1-r-a)-f\}. \tag{1}$$

The first month's new work has no same-month receipt. Month twelve's unpaid invoices remain receivables. Every failed request still consumes its modeled delivery cost; no expected invoice is created for it.

After startup and all incurred ordinary service/business costs, let cumulative operating income be Π_m^{op}. The tax-reserve target rises to $\tau\,[\Pi_m^{\mathrm{op}}]_+$ when that exceeds the existing reserve; no refund or tax benefit from discretionary construction is modeled. This is a planning reserve, not a jurisdictional tax calculation. Pass-through indirect taxes are outside net prices and can never become investment income. Recovered after-tax operating surplus is $S_m = [\Pi_m^{\mathrm{op}} - T_m]_+$.

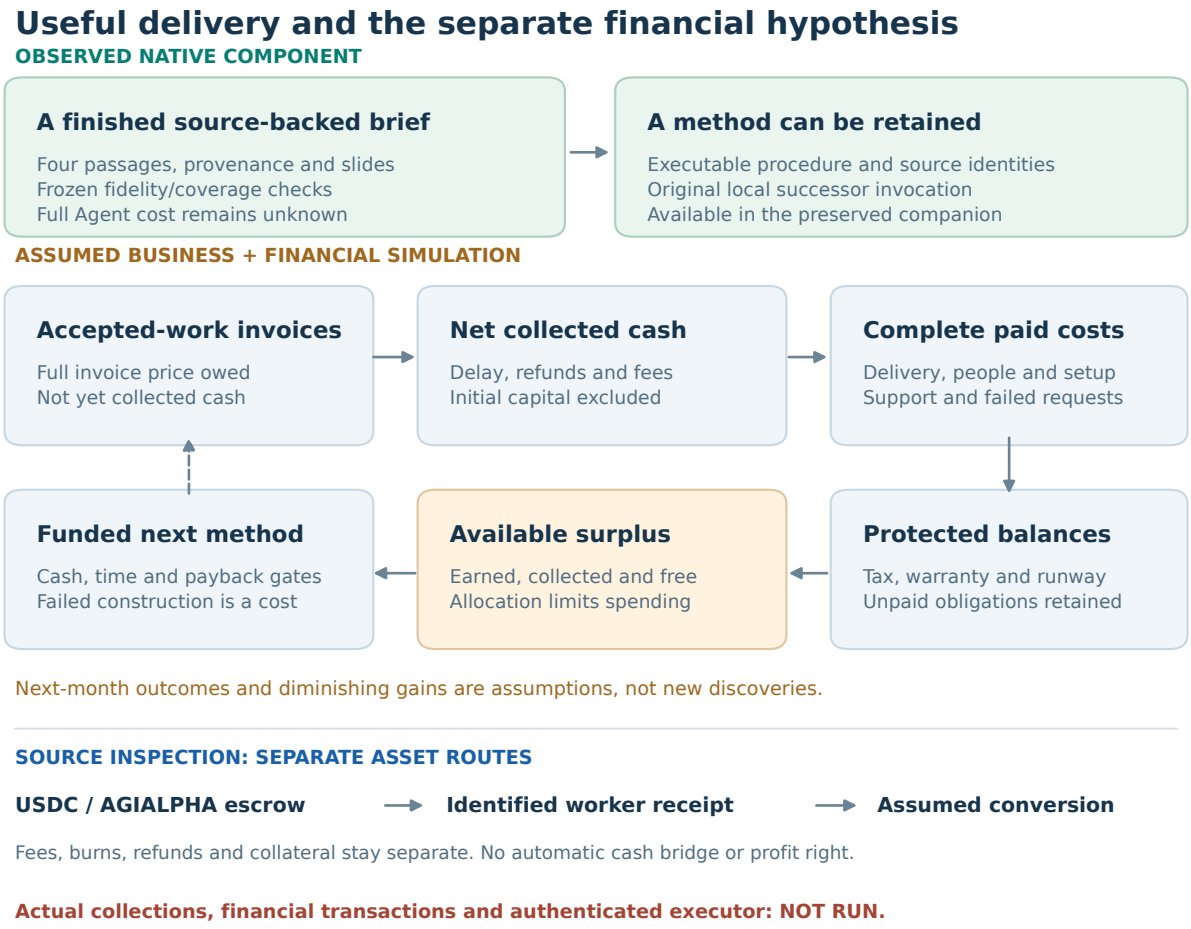


Figure 1. The economic engine separates the customer's alternative from the enterprise's cash. The new loop is a simulation of an explicitly proposed allocation policy. Existing native delivery and retained-module evidence keep their original scope; contract mechanisms do not create an automatic treasury rule.

Let C_m be remaining cash, W_m warranty reserve, K_m working capital, E_{m-1} current unspent earmarks and L_{m-1} lifetime allocations. For a chosen ceiling ρ, the new allocation is

$$\Delta E_m = \min\left\{[\rho S_m - L_{m-1}]_+, [C_m - T_m - W_m - K_m - E_{m-1}]_+\right\}. \tag{2}$$

This cumulative rule prevents repeatedly allocating the same retained surplus. Initial capital, uncollected invoices, collateral returns and asset appreciation never enter earned income. An allocation restricts cash; it becomes an expense only when spent. Final combined income deducts every discretionary construction and capacity purchase, including unsuccessful candidates. Mandatory pre-existing construction and upkeep remain in every policy, including zero new reinvestment.

3 Separate prospective specification

The 68-file specification was frozen on 25 October 2026 at 14:23:55.892155 UTC, after earlier native/operating results were known and before these financial calculations. Its public commit is `76ba78492da893b235b071cb0a61cf3608f9db37`; exact public readback and a coordination notice preceded the 14:25:47.798846 UTC GO. Independent prospective validation passed 282 static checks. Freeze SHA-256:

`4c7bac19106f946920d64ad39b01cb2cc844eed78c5cdb4d79140904a10cc9ba`

The primary matrix crosses two profiles, three joint operating scenarios, three face prices, four allocation fractions and two assumed capability-outcome worlds. All 144 paths are retained. The base USD20 cases add 0.25-times and four-times demand at every fraction/world, giving 32 further sensitivities. These are computational cases, not additional empirical task families or independent observations.[4]

The cash horizon is twelve operating months after setup. Initial virtual capital is USD25,000. Setup workdays are calculated from included paid provider hours and the declared daily resource allocation. This resolves the earlier cost paper's unvalidated setup-calendar shortcut without rewriting it. A startup requirement above virtual capital is refused before a partial operation is invented. Every scenario's imported service and improvement costs retain their original evidence class: modeled assumptions. The source brief's native receipt remains observed. Historical authoring costs remain unknown.[2]

The primary dollar labor basis is U.S. December 2025 office/administrative employer compensation of USD35.96 per hour; provider engineering/operation uses USD72.19 professional compensation. The preceding analysis also preserves Ontario and EU comparisons, source years, currency, benefits scope and task limitations. The wage sources do not establish each worker's AI use or this job's duration. Salary multiplied by a guessed time is not presented as a human experiment.[2, 5–7]

The hypothesis base needs only 0.10 provider minutes per expected accepted job to fit its USD0.50 variable delivery budget at the stated labor price. This is a demanding unvalidated requirement. Adverse hypotheses increase provider attention to one or five minutes per request and nonlabor cost twofold/fourfold, while acceptance, demand, collection and reserve conditions also deteriorate. The model therefore does not keep a fifty-cent cost while silently adding costly review. All paths charge required support and unsuccessful prior construction.

Table 1. Joint planning assumptions. These are not observed field rates or probabilities. Additional business costs supplement the imported complete service/improvement accounts; their stated labor is already included.

Parameter	Base	Conservative	Stress
Evidence-pack requested jobs/year	1,000	300	100
Hypothesis requested jobs/year	12,000	1,200	300
Acceptance probability	0.95	0.85	0.65
Collected fraction	0.98	0.90	0.75
Refund fraction of collections	0.01	0.05	0.10
Processing fraction	0.03	0.04	0.06
Processing USD per collected invoice	0.30	0.50	1.00
Additional business setup, USD	500	1,500	4,000
Additional monthly business cost, USD	150	300	600
Tax-reserve fraction	0.30	0.35	0.40
Warranty-reserve fraction	0.05	0.10	0.20
Working-capital months	2	3	4
Resource utilization	0.80	0.60	0.40

3.1 Resources, failure and complete demand

Requests are bounded by demand/backlog, provider minutes, customer minutes, executor occupancy and cash. Customer approval time cannot be removed by buying another executor. A shutdown incurs its full support obligation, pays available unrestricted cash and carries any unpaid balance. Later due collections repay that balance before free cash; service and investment do not restart.

For the evidence pack, the customer comparison includes its own preparation, review, correction and readiness. The complete-demand view also pays conventional workers for failed Agent requests and unserved demand, with fixed traditional enablement charged once. Customer prices are full face prices; a supplier's collection shortfall is not a legitimate saving for the customer. The USD200 hypothesis supplies replacement-cost arithmetic, but its human duration and fallback capacity remain unknown.

Resource limits are aggregate expected-use planning bounds. Provider token quotas, rate limits, network tails, daily staffing and live executor concurrency remain unvalidated. Fractional expected accepted outputs are not realized acceptances. The scenario reserve policy is not a probabilistic guarantee that real losses or liabilities fit the cash budget.

4 Customer and supplier economics

The evidence-pack base serves 1,000 requests, with 950 expected accepted outputs and 50 unresolved requests. At a USD20 accepted-job price, gross invoices are USD 19,000.00, but only USD 17,074.54 is collected within the horizon. Refunds and payment charges reduce that to USD 16,135.44. Provider startup, mandatory operations and all delivery attempts cost USD 18,765.44 before collection charges. Complete provider expense including processing is USD 19,533.80, or USD 20.56 per expected accepted output. Refunds remain a reduction in revenue.

The resulting provider income is negative USD 2,630.00; no earned surplus is allocated to new construction. Starting cash is a capital buffer, not profit. This USD25,000 account belongs to the provider and excludes customer outlays; it differs from the preceding paper's USD25,000 combined resource ceiling. The two admission quantities therefore must not be compared as an unexplained productivity improvement.

Table 2. USD20, zero new reinvestment, twelve operating months. Expense per accepted output includes payment processing and all incurred provider costs; income is after the modeled tax reserve. Zero-service plans have unavailable unit costs. All entries are simulations.

Profile	Case	Served	Accepted	USD/unit	Income USD
Evidence pack	Base	1,000	950.00	20.56	-2,630.00
Evidence pack	Conservative	125	106.25	245.36	-24,252.97
Evidence pack	Stress	0	0.00	N/A	0.00
Fifty-cent hypothesis	Base	12,000	11,400.00	2.61	121,120.44
Fifty-cent hypothesis	Conservative	1,200	1,020.00	37.66	-22,423.15
Fifty-cent hypothesis	Stress	0	0.00	N/A	0.00

For the assumed thirty-minute traditional production task, the full-demand customer cost is USD 32.27 per completed request including conventional rescue, against USD 28.42 for traditional work. The customer and provider both lose in this example. At an assumed sixty-minute traditional production duration, the customer instead compares USD 33.17 with USD 46.40; the provider still loses at the USD20 selling price.

The cohort-conditional provider break-even price is USD 23.21. At sixty assumed human minutes, the customer's upper price is USD 33.93. A price between those bounds can satisfy both modeled inequalities for the fixed cohort. It is a testable pricing interval, not observed willingness to pay or a rerun proving the same demand/admission path after repricing.

Table 3. Primary human-work sensitivities for the evidence-pack base cohort. Human durations are assumptions. The last column is a conditional interval between supplier break-even and customer indifference; endpoints are not commercial quotes.

Human minutes	Traditional USD/job	Customer price cap	Joint interval USD
5	13.43	0.97	None
15	19.43	6.96	None
30	28.42	15.95	None
60	46.40	33.93	23.21–33.93
120	82.36	69.89	23.21–69.89

Changing the human duration while holding the Agent case fixed does not establish that a more difficult job can be automated at the same cost. The actual four-passage brief has no matched human timing or customer-acceptance observation. All five duration values and all three selling prices remain in the raw outputs.

4.1 The USD0.50 / USD200 hypothesis

The base hypothesis requests 12,000 jobs, serves them within the declared aggregate model limits, and yields 11,400 expected accepted outputs. Setup is 48 modeled workdays before the operating year, with USD 5,703.30 provider/business startup cost. All of this is assumed operational behavior.

Table 4. Base fifty-cent hypothesis at USD20 and zero new investment. Collected money, costs and reservations are separate. Initial USD25,000 capital is excluded from income. These are not actual company receipts or expenses.

Account	USD
Accepted invoice face value	228,000.00
Gross cash collected	204,820.00
Refunds (contra-revenue)	2,048.20
Payment processing	9,216.90
Startup, operations and delivery before processing	20,525.70
Ordinary income before tax reserve	173,029.20
Tax reserve	51,908.76
Recovered income after tax reserve	121,120.44
Final uncollected invoice face value	19,000.00

Fifty cents is only variable delivery cost in this hypothetical offering. Full provider expense is USD 2.61 per expected accepted output after startup, service/improvement upkeep, business overhead and processing. The customer pays USD 20.36 including its own modeled attention per Agent-accepted output. Completing all requests, including 600 conventional rescues, costs the customer USD 29.49 per requested completion against the assumed USD200 alternative. That is a conditional 6.78-fold customer cost ratio, with no observed USD200 task, willingness-to-pay record or human fallback capacity. Its large modeled savings are not the provider's revenue.

The income after the tax reserve is USD 121,120.44. Warranty and terminal working-capital targets additionally restrict USD 1,862.00 and USD 1,520.40 of cash. Remaining free cash includes the original capital and must not all be called newly earned surplus. This favorable scenario establishes a conditional financial possibility; actual quality, attention, demand, collection and delivery cost would have to be measured for an offering with equivalent work.

The results require two separate adoption tests. The supplier must collect enough to cover complete costs and obligations; the customer must receive equivalent accepted work at a favorable full price including its time and fallback. A profitable supplier price can fail the customer test. Conversely, substantial customer savings can coexist with supplier losses. Neither a wage denominator nor a selected simulated margin establishes demand at the modeled price.

5 Bounded construction and the point at which to stop

Only previously recovered, collected and unreserved surplus can fund new construction. Three candidates are specified before confirmation. They have fixed construction/evaluation costs, included provider hours, upkeep, forecast benefits and simulated qualification outcomes. The first candidate fails in both outcome worlds; the no-gain world rejects every constructed candidate. Later hypothetical successes offer smaller fractional savings and greater construction costs. No result of this financial model is a discovered algorithm.

The investment gate forecasts savings on remaining demand under current resource constraints, current methods and current instances. It credits no future unfunded improvement. The current cash must already exist, and construction must fit unused provider time. Benefits begin only after

the next month's simulated qualification. A failure keeps its full cost. Negative remaining-horizon payback stops search; insufficient cash or time waits. A permitted allocation can therefore remain idle.

All 144 primary financial paths make zero new candidate or executor purchases. This is an operating result of the declared policy, not a claim that all future R&D is uneconomic. In the base hypothesis, the first candidate has a positive remaining-horizon forecast before any earned cash has arrived. By the time earned earmarks exist, its remaining reuse volume is below the frozen payback threshold. The twelve-month horizon deliberately provides no terminal credit for possible later benefits.

Table 5. Primary fifty-cent hypothesis: a percentage ceiling does not mandate spending. All four policies have the same operating income and accepted output.

Ceiling	Lifetime allocation USD	Actual new spending USD	Idle earmark USD
0%	0.00	0.00	0.00
25%	30,280.11	0.00	30,280.11
50%	60,560.22	0.00	60,560.22
75%	90,840.33	0.00	90,840.33

The fourfold-demand sensitivity demonstrates the conditional construction path. It serves 44,640 of 48,000 requested jobs; customer attention binds. At a 50% ceiling, USD500 is spent in month two on a candidate that fails its stipulated month-three qualification. A distinct USD750 candidate is built in month three and receives its assumed passing evaluation in month four. In the diminishing-gain world it is then used 33,480 times, saving USD 3,325.20 before its USD 135.00 upkeep and construction cost. The two candidates together yield USD 1,940.20 net simulated return before the tax-reserve treatment. Its gain is an input assumption, not a newly observed invention.

Table 6. Fourfold-demand sensitivity, fifty-cent profile and USD20 price. The accepted count is identical: the stipulated method improvement reduces cost but does not remove the binding customer-attention limit. All figures remain conditional simulations.

Policy/outcome world	Accepted	Spending USD	Income after reserve USD
Zero new investment	42,408	0.00	478,796.17
Half allocation; all candidates fail	42,408	1,250.00	477,546.17
Half allocation; stipulated diminishing gains	42,408	1,250.00	479,779.31

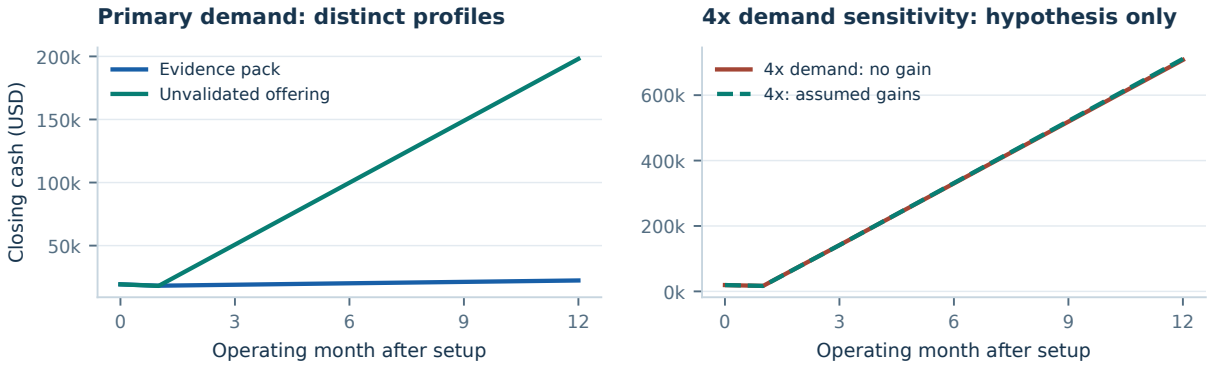
The successful candidate's realized-horizon payback requires 8,911 request uses, including its full horizon upkeep. Its probability-weighted forecast required 17,822 uses. The third candidate would require 154,624 forecast uses against 29,760 remaining resource-feasible uses, so construction stops. In the all-failed world the same funded attempts lose their full USD1,250. Neither positive demand nor abundant earmarked funds guarantees a useful next method. No executor is purchased in any reported path.

Buying execution capacity increases the available execution resource. Improving a method changes the cost or time of future work. Improving the discovery process would change the rate at which verified methods are found after complete costs. These are three different claims. The current simulation can test the financial conditions of the first two; it establishes none of them as a newly observed capability gain or an accelerating rate of improvement.

Simulated cash can remain idle

Base assumptions · USD 20 accepted-job price · 50% maximum surplus allocation

ALL FINANCIAL VALUES ARE SIMULATED. The USD 0.50 / USD 200 offering is unvalidated.



Case	Spent (USD)	Failed	Promoted*	Idle allocation (USD)
Evidence pack	0	0	0	0
Unvalidated offering	0	0	0	60,560
4x demand: no gain	1,250	2	0	238,148
4x: assumed gains	1,250	1	1	239,265

Spent includes failed construction. Idle allocations are restricted cash, not additional investment.

* Promotions are assumed outcome-world results. The predeclared 4x demand is not observed demand.

Closing cash includes remaining initial capital. Cash, profit, customer savings and capability gain differ.

Figure 2. Audited financial simulations distinguish operating receipts, reserves, funds earmarked and money actually spent. Failed construction and idle allocations remain visible. Differences between outcome worlds are assumed financial sensitivities, not empirical evidence of self-improvement.

6 What the token and contracts actually do

The inspected public sources cover all four repositories. The five-contract Alpha graph comprises Access, NovaSeed, RiskOracle, JobMarket and MARK; the canonical AGIALPHA token is an external dependency. Source inspection identifies economic rules and authority, while inspected tests and historical deployment reports have distinct evidence statuses. No new contract test or current-chain validation was performed in this amendment.[8–12]

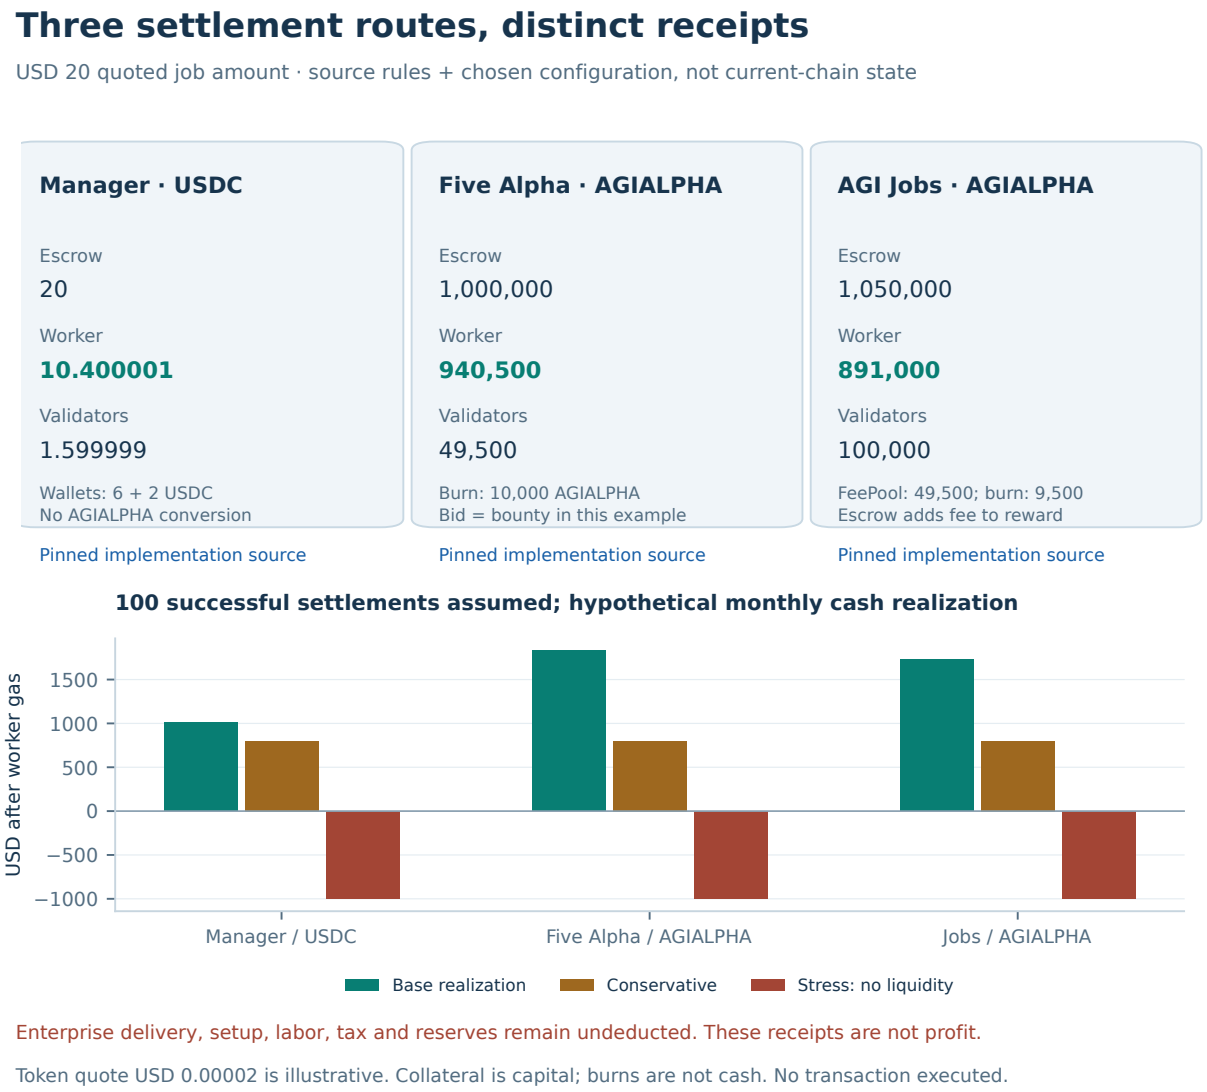


Figure 3. Distinct settlement mechanisms and recipients. Worker receipts are shown before enterprise delivery, support, taxes and reserves. Source formulas and configuration examples do not establish current live parameters or a reinvestment authorization. Currency-realization stresses are separate assumptions.

6.1 Manager USDC

For an evenly divisible accepted job with all reviewers correct, the inspected default reviewed-job split is 8% validators, 30% to the configured `wallet30`, 10% to `wallet10` and 52% to the agent. Integer residue and different settlement branches are retained in the executable formula. The two wallet labels do not establish ownership by the modeled enterprise, and the model credits only the worker's receipt. Returned performance bonds are capital. The worker bond includes a configured minimum, duration premium and caps; a flat percentage would omit implemented behavior. Expiry refunds do not deduct successful-job wallet shares.[13–15]

This is a USDC route. The inspected source establishes no automatic purchase, burn or appreciation of AGIALPHA from its activity. Generic release configuration retains deployment-required beneficiary fields; historical owner reports record their own paused-instance snapshots. Neither is current operational revenue evidence.[16, 17]

6.2 Ascension AGIALPHA

The client escrows bounty B and accepts a winning price $P \leq B$. A 5% gross reviewer pool is divided between two approving validators; each gross payout separately burns its rounded-down 1%. Without atomic rounding, the worker receives 94.05% of P, each approving validator 2.475%, and aggregate burn is 1%. Unspent bounty returns to the client. A worker locks the greater of 10% of bounty and the configured minimum stake. Default/rejection, missing quorum and losing bids have different collateral consequences.[11]

Those rules describe token quantities. Spending on dollar-denominated labor, compute or construction requires usable payment liquidity. The scenario's quote ratio is an illustrative accounting value; no current token market price or appreciation is assumed. Unconverted units remain inventory. A burn supplies no cash for the next model call.

6.3 Jobs v0, MARK and participant rights

Jobs v0 distinguishes quoted reward R from the additional job fee, worker payout weighting, worker-payout burn and FeePool-fee burn. A worker boost draws from a pre-funded operator reward pool. The shipped 5% fee, 10% validator reward and 1% burns are configuration examples with placeholder deployment fields, not independently verified current rates. Its source failure path can redistribute the employer's escrow under penalty rules; a universal full refund would be incorrect. Failed-route numerical economics are explicitly **NOT MODELED**.[18–22]

The 50/50 employer/treasury example concerns penalties, not a 50% appropriation of operating profit. FeePool rights belong to eligible configured stakers; they are not a dividend to every passive token holder. MARK routes pre-funded reserve only through its immutable-plan and proof constraints, with closing claims on unused contributed reserve. This does not distribute unrelated corporate profits. A company treasury policy would require a separately authorized appropriation; none is installed by this study.[12, 21, 22]

Table 7. Separate source-mechanism examples at quoted USD20, using the frozen illustrative valuation. The last column assumes 100 successful settlements, declared conversion charges/liquidity and gas allowances; it precedes all enterprise delivery/business costs, taxes and reserves. Jobs v0's USD20 is reward, so its customer escrow is USD21. None of these receipts is profit.

Route	Buyer escrow/job	Worker award/job	Monthly cash after gas
Manager USDC	20.00	10.40	1,012.92
Ascension AGIALPHA	20.00	18.81	1,827.88
Jobs v0 AGIALPHA	21.00	17.82	1,730.36

The conservative realization example caps converted cash at USD1,000 and pays USD200 gas allowance, leaving USD800 before business costs in each quoted-USD20 route. The stress example assumes zero conversion liquidity and USD1,000 gas allowance; cash after gas is negative USD1,000 despite token-denominated awards. These deliberately joint assumptions are not forecasts of token price or liquidity. Unconverted units remain inventory; collateral principal is not income; capital carrying cost is separately imputed.

Connecting a token settlement to an actual enterprise budget would require authorized recipient control, verified configuration, collection/conversion, collateral and refund accounting, and an

approved appropriation. The dynamic ordinary-USD study does not claim that connection has been implemented. The source modules expose its distinct payment boundaries without creating a new treasury or changing a deployed contract.

The historical Alpha report records five explorer matches and a scoped authority snapshot while withholding activation/public funding. Manager and Jobs reports likewise retain historical identities and paused-state boundaries. Fleet's in-process chain and mock-USDC profile are testing evidence, not public-chain paid operation. The source register preserves these distinctions and exact hashes. No source label is promoted into a current-chain verification.[17, 23–25]

7 Scale, launch scope and the next measurement

The practical unit is accepted work that a customer is willing to buy, supplied within verified resource and attention limits. The trillion-dollar labor context is separate. The preceding edition retains the U.S. USD5–10 trillion scenario based on compensation and assumed task shares, and treats the USD40 trillion global figure as illustrative. None of the present cash paths multiplies a selected task's cost advantage across that market.[1]

The immediate product can expose its bounded useful outputs, acceptance checks, limitations and retained-module behavior while broader capability research continues. A reliable commercial claim additionally needs matched traditional task observations or validated break-even thresholds, actual customer acceptance, observed willingness to pay, collection and support costs, and an approved live executor. These are specific measurements that reduce uncertainty. Recruiting a new worker panel is not required to report the current conditional analysis.

The next investment should target the binding economic uncertainty or delivery cost, then be tested on fresh useful work. Customer attention, collection, fixed-cost recovery and sufficient reuse can matter more than another reduction in an already small native computation time. Physical scientific or industrial applications also require domain validation, capital, energy, materials, manufacturing and lead times. Digital replication alone does not remove those constraints.

8 Reproduction and evidence status

The independent audit passes 1,288,965 arithmetic, ledger, identity and export comparisons, with zero remaining discrepancies. These counts are not independent observations or reliability trials. The scope includes all 176 paths, 2,112 monthly rows, 82,191 cost-ledger rows, 27 token realization cases and nine failure examples. A separate Decimal implementation reproduces the cash paths; token waterfalls use exact atomic integers. The 68 frozen files remain unchanged.

The first producer run is preserved. Independent review found 424 disagreements in two reporting boundaries, while all actual cash, delivery, cost and investment paths agreed. A documented *post-result reporting revision* makes stopped operational forecasts zero, labels optional counterfactual capacity separately, marks empty served-cohort comparisons unavailable and exposes complete provider expense including processing. All earlier values and attempts remain available; the revision changes no frozen assumption or actual financial path. Independent invariance checks and corrected export adapters are recorded rather than erased.

The model and result identities are bound in `evaluation/INDEPENDENT_AUDIT.json`. Cash results:

`e194f26361e69549b0e29181f2716e4229c63122eecc9b7c72307b4f48dc44eb`

Token results:

`aca12d6d1194f37967fda6fb2e831739b05571df1fd4a095f018c5f16f7f134a`

`MODEL_GUIDE.md` gives exact producer and safe independent-reproduction commands. The separate model never imports the producer's arithmetic. The source archive includes the monthly, cost, event and customer ledgers; prospective protocol; source snapshots; failed/declined candidates; and development records. No frozen destination is overwritten to reproduce the study.

The companion opens with the actual accepted brief and presentation. Native local work, recorded evidence and this financial simulation have explicit labels. The cash/token-flow diagram is editable; protocols, source snapshots, raw monthly rows, cost/event ledgers, code and independent checks are downloadable. The ordinary-USD model sends no network request and cannot execute a financial transaction. Authenticated browser-to-approved-executor and live settlement remain **NOT RUN**.

This amendment is additive. The original Economics edition, 42-job portfolio, negative method findings, frozen operating-cost analysis and earlier release boundaries are preserved. Financial feasibility can motivate further construction; the evidence from that construction must determine which capabilities are kept.

References

[1] AGI ALPHA research engineering study. AGI ALPHA FLYWHEEL: From useful work to reusable productive capability, 2026. URL https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/2afc163e4ca51ad85ae677b5fdffa3c69e2a662b/studies/agi-alpha-capability-construction-v1/README.md. Preserved 27-page edition and separately frozen evidence, 25 October; source commit 2afc163e4ca51ad85ae677b5fdffa3c69e2a662b.

[2] AGI ALPHA research engineering study. Complete economics of accepted work: separately frozen operational amendment. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/5caa871db78352557699f783027cf64357403d8b/studies/agi-alpha-operational-economics-v1/README.md, 2026. Canonical 14-page paper; conditional economics; native evidence preserved.

[3] National Center for O*NET Development. Secretaries and administrative assistants, except legal, medical, and executive: 43-6014.00, 2026. URL https://www.onetonline.org/link/summary/43-6014.00. Occupational task descriptions; retrieved 25 October 2026 for descriptive task mapping after the economic freeze; not a timing or quality measurement.

[4] AGI ALPHA research engineering study. Prospective collected-surplus and reinvestment protocol. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/76ba78492da893b235b071cb0a61cf3608f9db37/studies/agi-alpha-economic-engine-v1/FREEZE.json, 2026. 68 files frozen before financial calculations, 25 October 2026.

[5] U.S. Bureau of Labor Statistics. Employer costs for employee compensation—december 2025, 2026. URL https://www.bls.gov/news.release/archives/ecec_03202026.htm. Published 20 March 2026; Table 4, private-industry occupational groups; retrieved 25 October 2026.

[6] Employment and Social Development Canada. Technical writer wages in ontario, 2025. URL https://www.jobbank.gc.ca/marketreport/wages-occupation/20923/ON. Job Bank; Statistics Canada Labour Force Survey, 2023–2024 reference period; updated 19 November 2025; retrieved 25 October 2026.

[7] Eurostat. Eu hourly labour costs ranged from 12 to 57 euros in 2025, 2026. URL https://ec.europa.eu/eurostat/en/web/products-eurostat-news/w/ddn-20260331-2. 31 March 2026; covered enterprises with 10 or more employees; retrieved 25 October 2026.

[8] MontrealAI. MontrealAI/AGI-Alpha-Agent-v0: contracts/ascension/AscensionAccess.sol. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/f836f06e8351e64db1da770e16e304b5feb01f28/contracts/ascension/AscensionAccess.sol, 2026. Commit f836f06e8351e64db1da770e16e304b5feb01f28; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[9] MontrealAI. MontrealAI/AGI-Alpha-Agent-v0: contracts/ascension/NovaSeed.sol. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/f836f06e8351e64db1da770e16e304b5feb01f28/contracts/ascension/NovaSeed.sol, 2026. Commit f836f06e8351e64db1da770e16e304b5feb01f28; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[10] MontrealAI. MontrealAI/AGI-Alpha-Agent-v0: contracts/ascension/AscensionRiskOracle.sol. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/f836f06e8351e64db1da770e16e304b5feb01f28/contracts/ascension/AscensionRiskOracle.sol, 2026. Commit f836f06e8351e64db1da770e16e304b5feb01f28; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[11] MontrealAI. MontrealAI/AGI-Alpha-Agent-v0: contracts/ascension/AscensionJobMarket.sol. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/f836f06e8351e64db1da770e16e304b5feb01f28/contracts/ascension/AscensionJobMarket.sol, 2026. Commit f836f06e8351e64db1da770e16e304b5feb01f28; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[12] MontrealAI. MontrealAI/AGI-Alpha-Agent-v0: contracts/ascension/AscensionMark.sol. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/f836f06e8351e64db1da770e16e304b5feb01f28/contracts/ascension/AscensionMark.sol, 2026. Commit f836f06e8351e64db1da770e16e304b5feb01f28; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[13] MontrealAI. MontrealAI/AGIJobManager: contracts/AGIJobManager.sol. https://github.com/MontrealAI/AGIJobManager/blob/5d5670df9dbbc79da115e0a44b71989f34df58c1/contracts/AGIJobManager.sol, 2026. Commit 5d5670df9dbbc79da115e0a44b71989f34df58c1; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[14] MontrealAI. MontrealAI/AGIJobManager: contracts/utils/JobSettlement.sol. https://github.com/MontrealAI/AGIJobManager/blob/5d5670df9dbbc79da115e0a44b71989f34df58c1/contracts/utils/JobSettlement.sol, 2026. Commit 5d5670df9dbbc79da115e0a44b71989f34df58c1; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[15] MontrealAI. MontrealAI/AGIJobManager: contracts/utils/BondMath.sol. https://github.com/MontrealAI/AGIJobManager/blob/5d5670df9dbbc79da115e0a44b71989f34df58c1/contracts/utils/BondMath.sol, 2026. Commit 5d5670df9dbbc79da115e0a44b71989f34df58c1; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[16] MontrealAI. MontrealAI/AGIJobManager: config/usdc-deployment.json. https://github.com/MontrealAI/AGIJobManager/blob/5d5670df9dbbc79da115e0a44b71989f34df58c1/config/usdc-deployment.json, 2026. Commit 5d5670df9dbbc79da115e0a44b71989f34df58c1; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[17] MontrealAI. MontrealAI/AGIJobManager: config/mainnet/manager-8815-2026-10-17.json. https://github.com/MontrealAI/AGIJobManager/blob/5d5670df9dbbc79da115e0a44b71989f34df58c1/config/mainnet/manager-8815-2026-10-17.json, 2026. Commit 5d5670df9dbbc79da115e0a44b71989f34df58c1; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[18] MontrealAI. MontrealAI/AGIJobsv0: contracts/v2/implementation/JobRegistryBase.sol. https://github.com/MontrealAI/AGIJobsv0/blob/bc4abbd5c28f3672275b13300e40ea26c820b8ee/contracts/v2/implementation/JobRegistryBase.sol, 2026. Commit bc4abbd5c28f3672275b13300e40ea26c820b8ee; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[19] MontrealAI. MontrealAI/AGIJobsv0: contracts/v2/implementation/StakeManagerBase.sol. https://github.com/MontrealAI/AGIJobsv0/blob/bc4abbd5c28f3672275b13300e40ea26c820b8ee/contracts/v2/implementation/StakeManagerBase.sol, 2026. Commit bc4abbd5c28f3672275b13300e40ea26c820b8ee; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[20] MontrealAI. MontrealAI/AGIJobsv0: config/job-registry.json. https://github.com/MontrealAI/AGIJobsv0/blob/bc4abbd5c28f3672275b13300e40ea26c820b8ee/config/job-registry.json, 2026. Commit bc4abbd5c28f3672275b13300e40ea26c820b8ee; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[21] MontrealAI. MontrealAI/AGIJobsv0: config/stake-manager.json. https://github.com/MontrealAI/AGIJobsv0/blob/bc4abbd5c28f3672275b13300e40ea26c820b8ee/config/stake-manager.json, 2026. Commit bc4abbd5c28f3672275b13300e40ea26c820b8ee; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[22] MontrealAI. MontrealAI/AGIJobsv0: contracts/v2/FeePool.sol. https://github.com/MontrealAI/AGIJobsv0/blob/bc4abbd5c28f3672275b13300e40ea26c820b8ee/contracts/v2/FeePool.sol, 2026. Commit bc4abbd5c28f3672275b13300e40ea26c820b8ee; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[23] MontrealAI. MontrealAI/AGI-Alpha-Agent-v0: docs/agent/CURRENT_MAINNET_DEPLOYMENT.md. https://github.com/MontrealAI/AGI-Alpha-Agent-v0/blob/f836f06e8351e64db1da770e16e304b5feb01f28/docs/agent/CURRENT_MAINNET_DEPLOYMENT.md, 2026. Commit f836f06e8351e64db1da770e16e304b5feb01f28; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[24] MontrealAI. MontrealAI/AGIJobsv0: docs/production/mainnet-2026-10-18/deployment-snapshot.json. https://github.com/MontrealAI/AGIJobsv0/blob/bc4abbd5c28f3672275b13300e40ea26c820b8ee/docs/production/mainnet-2026-10-18/deployment-snapshot.json, 2026. Commit bc4abbd5c28f3672275b13300e40ea26c820b8ee; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.

[25] MontrealAI. MontrealAI/AGI-Fleet-Linux: docs/CONTRACT_PROFILES.md. https://github.com/MontrealAI/AGI-Fleet-Linux/blob/45759fb096faf8988f58ae5ea71248d0f35a97c7/docs/CONTRACT_PROFILES.md, 2026. Commit 45759fb096faf8988f58ae5ea71248d0f35a97c7; read-only source inspection 2026-10-25. No fresh contract execution or current-chain verification.