THE EMPLOYER'S DECISION: ACCEPTED WORK AT COMPLETE COST

AGI ALPHA FLYWHEEL

Complete Economics of Accepted Work

SUCCESSOR Ω HARNESS

Powered by TRANSFORMATIONAL INTELLIGENCE ∞

A separately frozen operational-cost amendment
to the preserved capability-construction study

Market-data-grounded comparison with competent traditional non-AI work.
Prepared with AI assistance for scientific and repository review.

Abstract

A useful delivery must be priced beyond the seconds spent computing it. We combine a measured native evidence-pack component with official labor prices, dated model and infrastructure tariffs, and explicitly assumed operational paths. Six mutually exclusive paths account for successful first delivery, repair, infrastructure retry, combined rescue, final rejection and early refusal. Three prospectively specified joint scenarios increase failures, review effort and resource use together. Finite demand pays for setup, approved-executor integration, deployment, unsuccessful capability construction, support and upkeep. We report service and improvement accounts separately before combining them, and add a complete-demand comparison in which conventional workers rescue unresolved jobs. In the base U.S. office-work case, the declared 1,000-request allocation costs $28.91 per expected accepted Agent output, versus $28.42 for an assumed 30-minute traditional production task with its ancillary work. That volume exceeds the example $25,000 resource budget. At the aggregate expected-use admission limits, costs are $31.74 for Agent-accepted work, $28.52 for traditional completion and $35.54 for complete-demand human fallback. These are planning ceilings with unvalidated daily scheduling, not hard stochastic budget guarantees. Conservative and stress cases cannot fund their fixed resource commitments within the stated budget. These are conditional deployment economics, not measured invoices, recruited-human results or proof of operational readiness. The earlier 42-job study and both original research freezes remain unchanged. Authenticated browser-to-approved- executor testing remains **NOT RUN**.

Measured native work + externally sourced prices + declared operational assumptions.
No simulated provider call, human review or deployment is reported as observed.

1 Start with a finished deliverable

A project coordinator needs a compact evidence pack before publishing a research paper and planning executor integration. The brief supplies two approved, pinned documents and four specific passages. The requested result is a readable cited brief, a short presentation, source identities and a clear statement of limits. This is useful administrative and technical-document work: it collects the exact publication route and executor prerequisites in a portable, inspectable artifact. It is deliberately modest. Four supplied passages do not establish open-ended research, semantic synthesis or a general professional role. O*NET describes report/presentation preparation and information search among administrative assistant tasks; this is a task mapping, not a measure of duration or of the share of that occupation automated.[1]

The preserved native job `source-synthesis-01` produced the HTML/JSON brief and presentation, source ledger, lineage and retained executable module. Its measured delivery boundary, including the local acceptance call, was 0.111703 seconds; child plus supervisor/oracle CPU was 0.124822 seconds. It made zero task-time model calls. Employer attention and complete monetary cost were unmeasured. Separately implemented deterministic acceptance checked exact required passages, source binding, artifact hashes and usable structure. It did not establish the truth of the source documents or independent human acceptance.[2]

The source texts and requested passages are the same for both economic arms. A competent traditional worker may use ordinary non-AI software, templates, copy/paste, search and professional judgment efficiently. No guessed duration is presented as a timed worker. The operational Agent scenario retains the native component but assumes additional intake, orchestration, approval, review and exception handling. Its proposed model calls were neither necessary for the frozen exact-passage computation nor executed in this study.

Table 1. The transaction being priced. Both workflows must meet the same stated artifact and review standard; durations outside the native component are assumed.

Requirement	Traditional workflow	Agent service scenario
Intake	Read the supplied brief, passages and source access; clarify scope.	Validate a bounded request; retain human scope and approval responsibility.
Produce	Assemble and format the cited brief/slides using ordinary tools.	Invoke the actual local producer; assumed provider calls handle orchestration and review assistance.
Verify	Check coverage, fidelity, usability and material objections; correct and recheck when required.	Run native checks; charge human review, verified challenges, retries, repairs and exhaustion handling.
Deliver	Supply readable files, source references and limits.	Supply those same outputs plus machine-readable evidence and the retained module.
Continue	Reuse templates and professional familiarity.	Import the actual retained module on a changed mission; preserve valid dependencies and fall back when the candidate offers no advantage.

The browser companion leads with these finished files and acceptance evidence. Its source-brief, changed-source repair and successor presentation run locally. Recorded native planning results and modeled operating economics carry separate labels. The retained module's exact bytes are verified before import. Retention demonstrates an actual dependency and transferable function; the preserved mechanism tests do not show a CPU advantage or self-acceleration.

2 A new analysis with a separate freeze

This amendment was specified after the native experiments were known and before its new scenario calculations. It is a prospective computational economic analysis of declared assumptions, not a retrospective expansion of empirical endpoints. The original Economics edition associated with #4901, its 69 frozen inputs, and the expanded study's 110 frozen inputs and all outcomes remain intact. The 16-file economic freeze was created at 13:29:21.657977 UTC on 25 October 2026. Its public source commit is `4d1d2af3e77a5346e3dc1a574f509d5c67c0aef7`; the public prospective notice preceded the 13:30:40 UTC authorization to calculate. Independent pre-calculation review passed 61 specification, identity and source checks. Manifest SHA-256:

`54bc70fec9bb635b473193d966d36ed5bf445460a58255f025cbed4a30854d8b`

The actual native outcomes were already known. The amendment therefore cannot serve as a preregistration of those outcomes, and it is not described as one.

Table 2. Evidence classes. A published price does not make assumed usage observed.

Class	Examples	Interpretation
Observed component	Native elapsed/CPU receipt, exact output, zero task-time model calls, original acceptance results.	Local boundary only; not total delivery time or an invoice.
External observation	Employer labor costs, occupational wages, API tariffs, infrastructure rates and reference exchange rates.	Dated price observations; do not establish workflow duration, quality or deployment behavior.
Operational assumption	Tokens, calls, failure paths, review effort, queue waits, construction, support, demand and utilization.	Joint scenarios and sensitivity ranges; not empirical probabilities or confidence intervals.
Unknown history	Research-authoring tokens/USD, human minutes, complete live deployment costs and customer acceptance.	Remain unknown. A simulated future budget is never substituted for the missing historical bill.

Official labor observations retain year, geography, occupation, currency, benefits coverage and methodology. The full companion preserves the original 24-record register across the United States, Canada and Europe and its O*NET task mappings. Wage-source AI use is not established. Published programmer and professional-writing experiments supply adjacent evidence, but their task scope, completion conditioning and AI-use limitations prevent importing their durations into this evidence-pack job.[3, 4] Recruiting a new panel is unnecessary for a transparent break-even calculation and remains outside this analysis.

3 Price the entire request, including failure

Let p_b be the probability of terminal path b, $a_b \in \{0, 1\}$ indicate a genuinely acceptable deliverable, and c_b include every charged event reached on that path. Expensive failed attempts are not removed from the numerator. With N requested jobs and finite-horizon fixed cost F,

$$Q = \sum_b p_b a_b, \qquad \bar{c} = \sum_b p_b c_b, \tag{1}$$

$$E[A] = NQ, \qquad C_A = \frac{F + N\bar{c}}{NQ}. \tag{2}$$

This ratio of expectations is a planning quantity, not the expected ratio of a small random sample. Zero demand or zero acceptance produces an undefined accepted-normalized cost, not zero-cost work.

The six paths are direct acceptance, model repair, infrastructure retry, combined retry/manual rescue, final rejection and early refusal. Their event sequences link failed attempts to additional tokens, operator diagnosis, human review, repair and waiting. At most three production attempts are allowed. This avoids the implausible assumption that hard requests become independent easy requests whenever a retry begins. Base, conservative and stress weights also change usage, demand and utilization together; they are not fitted to the 42 accepted fixtures.

The service account includes model/tool use, allocated executor occupancy, storage and operations, employer approval/review/correction, provider support, setup, integration and deployment. The improvement account includes useful and unsuccessful construction, independent qualification, regression checks, capability refresh and upkeep. Each component has one owner and charging rule. The combined comparison is attributed resource cost for self-operation or a cost-transparent service. It is not a provider quote. Adding an unspecified provider fee to the same provider costs would double count that business boundary.

Provider engineering and operation use the U.S. professional compensation rate, including for the Ontario and EU price illustrations. Customer and traditional-worker time use their stated local labor basis. These are explicit labor-cost proxies, not observed contractor fees.

The initial investment is paid once, recurring support and upkeep are paid for the declared horizon, and all requested jobs pay their path costs. We show the first commission separately, charging the initial investment and one operating month before any future reuse. A low unit cost at high assumed volume is not a low cash requirement for a new pilot. Historical research spending remains unknown and is not silently replaced by these forward construction budgets.

3.1 Published price basis and explicit planning budgets

The primary administrative-brief comparator is U.S. private-industry office and administrative compensation of $35.96 per hour in December 2025, comprising $24.79 in wages and $11.17 in benefits. Provider engineering/support uses professional compensation of $72.19, comprising $50.10 in wages and $22.09 in benefits. Neither is a contractor quote or a matched worker observation.[5] Ontario technical writers have a CAD 34.98 hourly median in the 2023–2024 Labour Force Survey, published through Job Bank; the 1.30 employer multiplier is assumed. The EU example uses EUR 34.90 covered-economy employer cost for 2025, which is not an occupation-matched wage.[6, 7]

All four price bases across three geographies retain their scope differences. USD technology/provider costs are nominally translated using 23 October 2026 ECB references: EUR 1 = USD 1.1193 and CAD 1.5665. These are reference conversions, not purchasing-power parity or local cloud-SKU invoices.[8] Mixed labor-price vintages remain visible; no undisclosed wage-growth adjustment turns them into 2026 wage observations.

The hypothetical first-party Sonnet 5.5 route costs $2 per million input tokens and $10 per million output tokens; Opus 5.5 costs $4 and $20. Successful search is $10 per thousand uses. Declared token totals include context, tool/history and generated reasoning. No cache, batch or free-tier discount is claimed.[9] Cloud Run request-billed Iowa rates are $0.000024 per allocated vCPU-second, $0.0000025 per GiB-second and $0.40 per million requests, with 100-ms rounding. The assumed executor has two separate one-vCPU, two-GiB instances with one request per instance. Standard regional storage is $0.000027397 per GiB-hour; operations/egress have explicit separate allowances.[10, 11] All prices were retrieved on 25 October 2026. No provider was called or deployed to validate the assumed usage, latency or quality.

Table 3. Joint operational assumptions. "Base" is an illustration, not an estimated most-likely case. Every value in this table is assumed.

Parameter	Base	Conservative	Stress
Requested jobs in annual illustration	1,000	300	100
Delivery token multiplier	1.0	2.0	4.0
Delivery human-effort multiplier	1.0	1.5	2.5
Executor occupancy multiplier	1.0	1.5	3.0
External-wait multiplier	1.0	2.0	4.0
Initial provider engineering hours	72	184	392
Recurring provider hours per month	8	18	36
Initial employer onboarding hours	2	4	8
Employer governance hours per month	1	2	4
Assumed setup workdays	10	30	60
Available-resource utilization	0.80	0.60	0.40

The initial base budget includes 16 hours of useful capability construction, 16 hours of unsuccessful construction and 8 hours of qualification, in addition to 8 hours of requirements, 8 of deployment and 16 of executor integration. Recurring base provider work is eight hours per month across support, runtime maintenance, recovery upkeep, capability checks, refresh and unsuccessful search. These are inspectable engineering allowances, not estimates of historical authoring time. The monetary resource budget is USD 25,000 in all primary examples. Employer and provider pools each reserve 120 minutes per workday; traditional production has a separate 480-minute worker pool. The year has 240 eight-hour workdays. Token pools and utilization are assumed planning resources, not verified provider quotas.

3.2 Time and attention have different denominators

Employer minutes include active intake, approval, review, correction and governance. Provider engineering and operator minutes are shown separately, then combined as total human labor. Successful-request turnaround conditions on acceptance: $\sum_b p_b a_b \ell_b / Q$, where ℓ_b is nonoverlapping elapsed time. The all-request mean, rejected-request mean and discrete median/p95 remain visible. Accumulated serial workload per accepted output is not a customer's latency. Executor occupancy contains model/network/native work; the measured native duration is inside it and is not added a second time. The native CPU reading is not a measurement of allocated cloud-instance time.

3.3 Human fallback completes the demand that Agent acceptance leaves open

An accepted-normalized cost can conceal unresolved requests. We therefore show a separate hybrid: every final rejection or refusal goes to a competent traditional worker for the complete workflow, plus an explicit handoff. It pays the entire failed Agent path, the full human rescue, both routes' relevant fixed costs and all review. No partial Agent work is credited as free human progress. Assumed final human acceptance is one for both the comparator and the rescue; this favors competent traditional work and is not an observed success rate.

4 Results: where the comparison changes sign

The primary base case does not establish a cost or attention win. At the example budget's aggregate expected-use limits, Agent-accepted work costs $31.74, traditional completion at assumed $t = 30$ productive minutes costs $28.52, and the complete-demand hybrid costs $35.54 per output. Their admitted cohorts differ because each pays its own fixed and variable costs. The Agent-only cohort still has unresolved requests; the hybrid explicitly finishes all requests it admits using assumed competent-human rescue.

Table 4. Base U.S. office-work example under a USD 25,000 economic budget. These are aggregate expected-use planning ceilings, not a validated schedule or a hard budget guarantee. Productive human time is assumed 30 minutes.

Route	Requested	Expected completed	Cost per accepted	Employer min/output
Agent-accepted	829	787.55	$31.74	19.51
Traditional, $t = 30$	876	876.00	$28.52	15.88
Agent + human fallback	703	703.00	$35.54	20.33

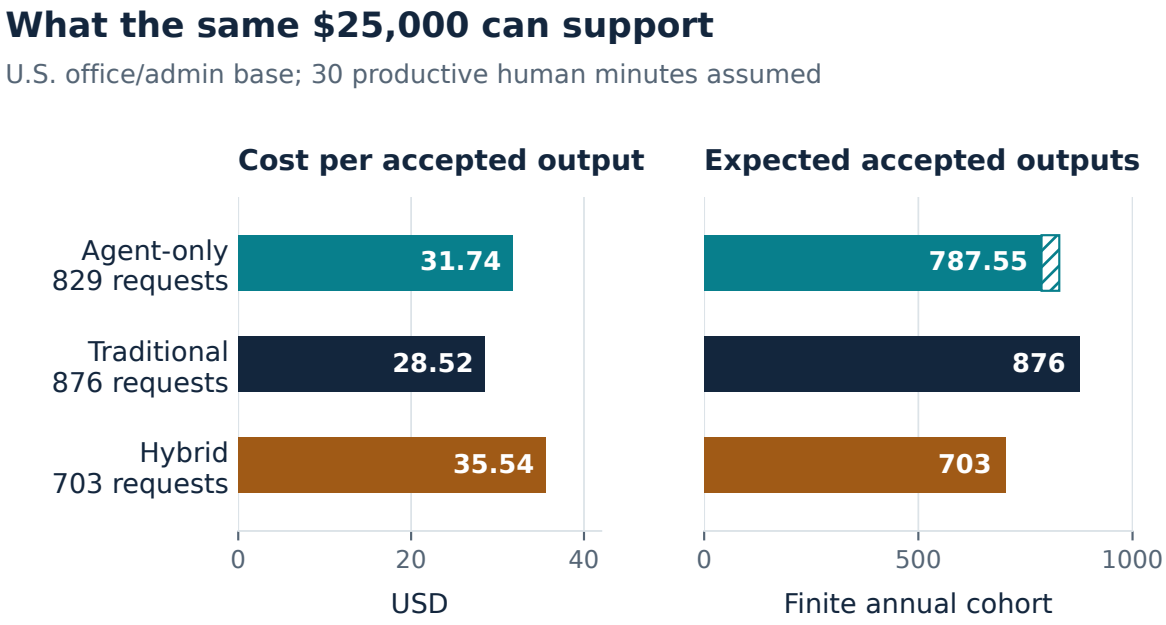


Figure 1. Complete cost and accepted capacity at the aggregate resource limits. The incomplete-demand Agent route and complete-demand fallback are different offerings. This is simulated service economics, not a live invoice.

The original 1,000-request allocation would cost USD 27,469.20 for the Agent, above the USD 25,000 budget, and would produce only 950 expected accepted outputs. Its accepted-normalized USD 28.91 is therefore not the attainable unit cost under that example budget. At the same declared volume, the traditional route costs USD 28.42 per completed job and the hybrid USD 29.73. All three fail

to fit the entire 1,000-request demand within the example budget. Cost and capacity must be read together.

Table 5. Complete U.S. office-work costs at declared annual demand. The human column assumes 30 productive minutes plus its ancillary work. The Agent columns are conditional calculations at volumes that are not feasible within the stated budget; conservative/stress fixed commitments alone block admission.

Scenario	Requests	Agent Q	Agent	Traditional	Hybrid
Base	1,000	95%	$28.91	$28.42	$29.73
Conservative	300	85%	$152.77	$38.68	$140.54
Stress	100	65%	$1,066.77	$76.12	$742.75

The base Agent's annual fixed account is $13,029.14. The full annual declared-volume account separates $20,718.70 of service/employer resources, $6,352.72 of identified improvement work and $397.78 of shared service/improvement resources. Shared costs remain in the total even when attribution between service and improvement is unresolved. This avoids making either account look cheaper through an arbitrary favorable allocation.

The first commission is a different decision. Its initial investment is $5,275.22, its one-month operating reserve is $646.16, and its probability-weighted request cost is $14.44. Combined, it costs $5,935.82 per requested commission, or $6,248.23 normalized by assumed acceptance. This is a modeled service-construction investment, not an invoice or a claim that a customer should pay thousands for four passages. A new service built for one brief is economically different from using an already available, frequently reused capability.

At declared annual volume, Agent employer attention is 19.33 minutes per accepted output, versus 15.82 for traditional work. Yet total human labor is 33.20 versus 47.10 minutes. Higher-priced provider effort and fixed construction explain how less total labor can coexist with greater employer attention and slightly greater cost. Neither released labor nor faster computation is automatically a payroll saving.

Successful Agent requests have an assumed mean turnaround of 40.88 minutes, median 34.17 and p95 55.17; traditional mean turnaround at $t = 30$ is 56.20 minutes, and hybrid mean turnaround is 45.17 minutes. These are event-path elapsed-time assumptions, not measurements of provider service. The native 0.111703-second component cannot replace those clocks.

Table 6. Base scenario price illustrations at the same declared 1,000-request allocation, which is budget-infeasible for all four price bases. Values are in each row's currency. Thresholds concern productive human minutes, not all-in human effort. No geography-specific performance inference is made.

Labor price basis	Currency	Agent	Human $t = 30$	t^* min
U.S. office/admin	USD	28.91	28.42	30.83
U.S. professional	USD	40.58	56.86	16.47
Ontario technical writer	CAD	38.90	35.96	33.88
EU covered economy	EUR	26.73	27.57	28.56

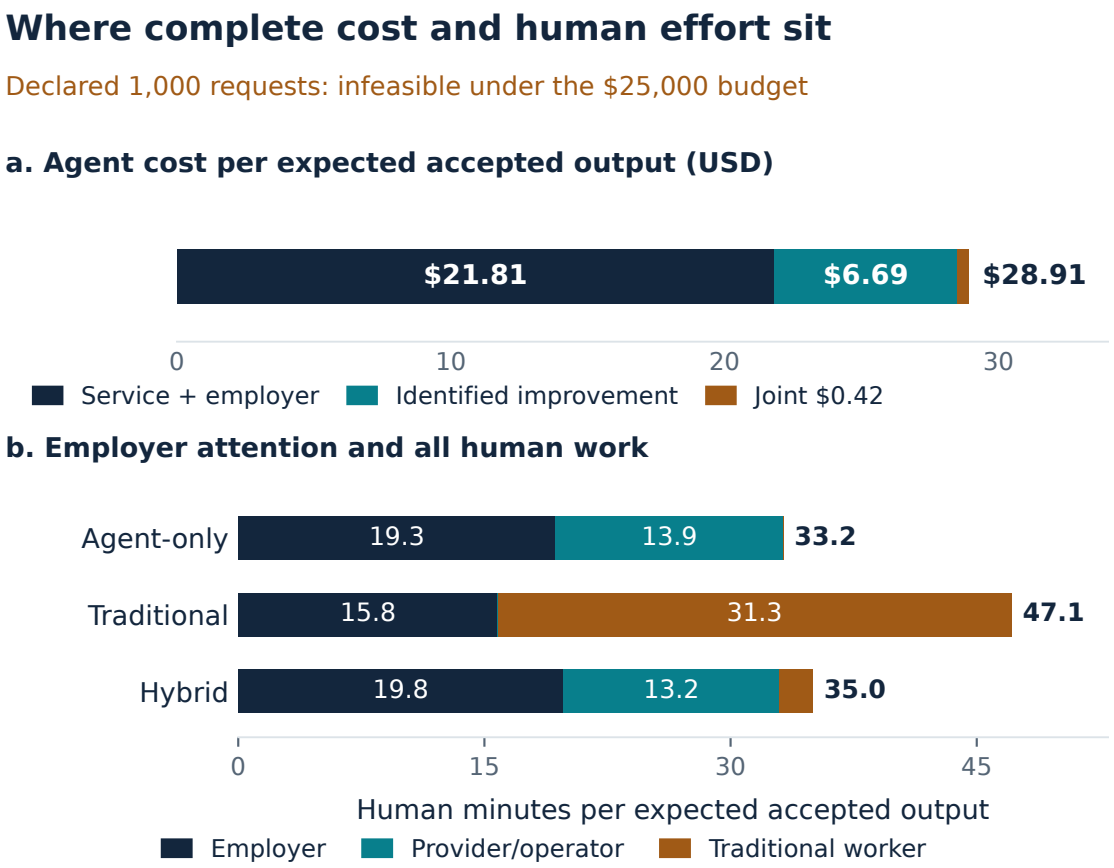


Figure 2. Money, remaining employer attention and total human labor answer different questions. These declared-volume accounts include fixed costs; the 1,000-request illustration exceeds the example resource budget.

5 Break-even replaces an invented worker

Traditional productive time t covers source inspection, assembly, formatting and the worker's own check. Client intake, independent review, correction and delivery are charged separately. Prespecified values are 5, 15, 30, 60 and 120 minutes. Five minutes is a material case for four supplied passages; larger values are sensitivities, not statements about how long this job takes. Ordinary software and a reusable-template setup allowance are explicit. Both arms can learn and reuse legitimate methods.

At fixed demand, traditional cost is $C_H(t) = H_0 + w_H t/60$, so the productive-time break-even is

$$t^* = \frac{60(C_A - H_0)}{w_H}. \tag{3}$$

H_0 contains traditional ancillary work and fixed allocation. The model also shows the corresponding all-in human minutes. A customer can replace the assumed productive duration with a timed conventional workflow of matching scope and acceptance. A negative threshold would mean the assumed Agent cost is below the traditional ancillary cost; it would not imply negative work time.

At the declared 1,000-request base allocation, the Agent-only productive-time threshold is 30.83 minutes. Complete-demand fallback needs 32.31 minutes. These thresholds preserve human ancillary and fixed costs; neither means the four-passage brief has been observed to require that long. The 1,000-request allocation remains outside the example budget.

Table 7. All one-way sensitivity families around the U.S. office base. The final column is declared-volume cost per expected accepted output in USD; feasibility must be read from the full output. These are assumption ranges, not confidence intervals or probabilities.

Driver	Frozen values	Cost range
Requested volume	1, 10, 100, 1000, 10000	16.57–13,730.08
Delivery tokens	0.5, 1, 3, 10, 30	28.83–33.70
Employer review effort	0.5, 1, 2, 4	26.59–42.86
unsuccessful initial construction human hours multiplier	0.5, 1, 2, 4	28.31–32.56
Recurring support/upkeep	0.5, 1, 2, 4	24.83–53.40
Native route; no delivery model/search	True	28.72–28.72

Removing all hypothetical delivery model calls and paid search leaves the declared-volume cost at USD 28.72. That sensitivity retains positive setup, integration, construction, review, occupancy and upkeep budgets. It is not evidence that the untested full service has identical quality without those calls. It shows that token prices are a small part of this assumed service's cost. The native route's actual zero task-time model calls do not erase the unmeasured cost of authoring and operating the capability.

A separately labeled post-result algebraic diagnostic solves for reuse volume using the same fixed/variable accounts, with no new parameter search. At the U.S. office $t = 30$ baseline, Agent-only parity requires 1,040 requests; complete- demand fallback requires 1,113. The corresponding aggregate admission limits are 829 and 703. Thus neither payback volume fits the configured resource plan. The formulas and independent substitution checks are supplied in `outputs/post_result_reuse_diagnostic.json`. This is not a retroactively preregistered primary endpoint. None of the four price bases' parity thresholds fits both arms' same-request-cohort resource bounds.

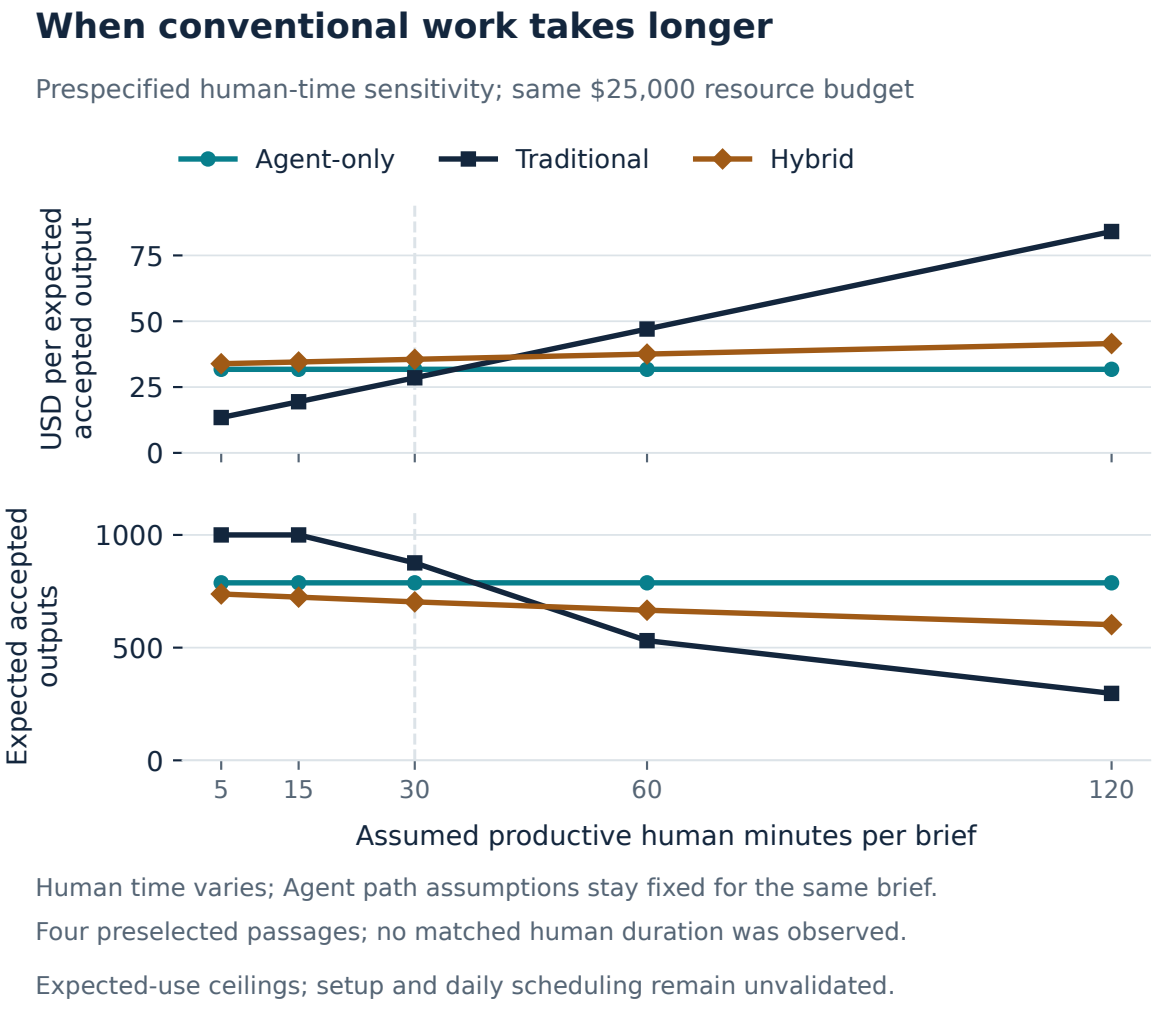


Figure 3. Prespecified human-duration sensitivities at each route's aggregate expected-use admission limit. Agent assumptions stay fixed while hypothetical human time changes. This does not demonstrate a harder task at constant Agent cost; no matched human duration has been observed.

The apparently attractive USD 16.57 cost at 10,000 requested uses is also outside the resource plan. Its admitted Agent volume remains 829 and its admitted unit cost USD 31.74. Unlimited amortization would therefore create an unsupported commercial claim. A real decision needs adequate demand and the money, human time and infrastructure to serve it.

6 Capacity and commercial interpretation

The model constrains feasible requests jointly by demand, a finite attributed economic budget, employer and provider minutes, executor occupancy, model-token budgets and utilization. It subtracts fixed work from the corresponding pools and removes setup days from the service calendar. Accepted capacity is N^*Q, where N^* is the integer number of requests admissible under every resource bound. Reported expectations may be fractional; they are not guaranteed integer service commitments. Costs at declared demand are labeled infeasible if the same scenario cannot serve that demand. Amortization over unattainable volume cannot establish a feasible low unit cost.

Table 8. Base U.S. office expected accepted capacity at USD 25,000, with each route constrained separately. Human time is an assumption. Agent admitted cost is USD 31.74 throughout this duration sensitivity. Capacity is not guaranteed.

Human productive min	Agent	Human	Hybrid	Human cost	Hybrid cost
5	787.55	1,000	738	13.43	33.87
15	787.55	1,000	724	19.43	34.53
30	787.55	876	703	28.52	35.54
60	787.55	531	666	47.04	37.53
120	787.55	297	602	84.08	41.52

Longer hypothetical conventional production can produce conditional gains. It does not establish that this small source pack takes 60 or 120 minutes, nor that Agent costs and acceptance would remain fixed on a more difficult job. At the 30-minute U.S. professional price basis, the separate-cohort aggregate comparison is 435.10 expected Agent-accepted outputs versus 426 traditional outputs, but only 362 hybrid completions. Those modest conditional differences do not support a broad claim of autonomous enterprise superiority.

The capacity equations do not validate a staffing schedule or a hard spending cap. A post-result consistency check found that the assumed Agent setup windows of 10, 30 and 60 workdays require at least 45, 160 and 490 provider workdays under uniform daily availability and utilization. Frontloading an annual-average labor pool is a different scheduling assumption, not an observed deployment. The traditional stress setup also fails its uniform employer-time limit (five days needed versus four declared). Frozen values remain unchanged and the diagnostic is published separately. Declared first-commission setup latency must therefore not be promised as achievable.

Moreover, admission is sized against probability-weighted resource use. Actual costly failure paths could exceed a budget that is nearly exhausted at its mean. No chance constraint, tail-risk reserve or paid-call spending enforcement is validated here. The original study's real local STOP/work-unit controls remain separate. Prospective paid operation would need its own measured, authorized resource guards and staffing validation.

Freed employee minutes have economic value when reassigned to useful work. They become payroll savings only when a specific employment or purchasing decision actually reduces expense. A cheaper computational path is not customer revenue, provider margin or cash available for reinvestment. This study reports no observed sales or reinvestment. The broader study's bounded

reinvestment simulations and U.S. $5–10 trillion digital-work scenario remain separate; neither scales this small job's acceptance or savings across the labor market.

Further construction should stop when its complete incremental cost exceeds expected savings across eligible, feasible future demand, or when it fails the quality/resource boundary. The preserved native candidate losses make that rule concrete: the available specialist remains selected where the retained method does not improve the result. The capability archive can retain useful functional methods without attributing a nonexistent speed advantage to them.

7 What the next measurement should resolve

The most useful next field measurements are matched-scope conventional task time and quality; active employer/provider minutes with correction events; per-request provider token/tool receipts; real retry and exhaustion paths; and paid demand, support and maintenance over a finite horizon. Measuring native milliseconds more precisely would do little to resolve those unknowns. The scenario driver analysis orders the practical uncertainty; its ranges are not calibrated probabilities.

For this low-consequence document task, the scenario cost boundary covers the declared delivery and operating resources. It does not invent actuarial losses for undetected consequential errors. Regulated clinical decisions, financial transactions, restricted data, physical actions and other unsupported categories remain outside the measured offering. A future use with material downstream liability needs its own acceptance and loss model.

The offline companion demonstrates fresh local computation, retained-module invocation and recoverable artifacts. It does not establish authentication, continuous service, deployment uptime or browser-to-approved-executor behavior. That specific integration remains **NOT RUN**. Completing this economic analysis does not remove the restriction or manufacture a live test.

8 Reproduction and release boundary

The independent arithmetic implementation reconciled 157,435 checks across all 12 core scenarios, 60 human-duration and corresponding hybrid comparisons, and 23 one-way sensitivities. A separate derivative audit checked the reuse equations and all 5,064 component-ledger rows. These are accounting checks, not additional empirical jobs or human observations. The first audit adapter's enum/reference-selection errors and their correction remain recorded; no producer arithmetic or frozen input was changed to clear them. A clean producer rerun reproduced all five primary output files byte for byte.

The full source contains the model, its independent Decimal implementation, the six event paths, price/source registers, observed native artifacts, frozen specification, CSV ledgers and all output JSON. From the extracted project root:

```
python economics_amendment_v1/src/model.py --help
python economics_amendment_v1/paper/build_content.py
make -C economics_amendment_v1/paper
```

`MODEL_GUIDE.md` gives the exact model/audit invocation and output directory safeguards. The preserved native study has its own reproduction instructions and unchanged complete LaTeX archive. Figures retain editable Matplotlib source and SVG/PDF/PNG exports. The browser's download panel includes the source-linked economic inputs, model and audited results.

Canonical scenario-result SHA-256:

`cba5f2cf284e27f7097fc3f4cebe0bf985952cd64bfcabee78af51ca0c17ac9d`

The separately reviewed PDF and final archive manifests bind the delivered bytes. A public draft is inspectable research material; peer review, repository merge, website publication and live operating admission remain separate.

The complete expanded portfolio remains 42 primary bounded jobs across seven families. Repair/transfer, review, controls and repeated load attempts retain their own accounts. They do not add independent occupational coverage. All failed construction candidates and unfavorable economic rows are preserved. No evidence in this amendment establishes self-acceleration.

The original Economics publication proceeds independently through #4902; its source record remains #4901. The preserved expanded paper is proposed in #4904, with research evidence in #4903. Prepared artifacts, public draft source, review, merge, website publication and operational admission are distinct states. The accompanying dated publication-status record identifies the actual state; this manuscript does not imply that a proposed paper has been merged or a service deployed. Existing owners, release boundaries and approval controls remain in force.

References

[1] 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.

[2] 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.

[3] Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The impact of ai on developer productivity: Evidence from github copilot, 2023. URL https://arxiv.org/abs/2302.06590v1. arXiv:2302.06590v1; task- and completion-conditioned control results are not a matched timing benchmark for this study.

[4] Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence, 2023. URL https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf. 2 March working-paper edition; 444 participants; control AI contamination precludes strict no-AI timing calibration here.

[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] European Central Bank. Euro foreign exchange reference rates, 2026. URL https://www.ecb.europa.eu/stats/shared/pdf/eurofxref.pdf. 23 October 2026 reference date; retrieved 25 October 2026; nominal reference rates, not purchasing-power parity.

[9] Anthropic. Pricing: Model tokens and tool use, 2026. URL https://platform.claude.com/docs/en/about-claude/pricing. First-party standard Claude Sonnet 5.5 and Opus 5.5 API prices and web-search charges; retrieved 25 October 2026.

[10] Google Cloud. Cloud run pricing, 2026. URL https://cloud.google.com/run/pricing. Services, request-based billing, default USD rates for Iowa us-central1; retrieved 25 October 2026.

[11] Google Cloud. Cloud storage pricing, 2026. URL https://cloud.google.com/storage/pricing. Standard regional at-rest storage, Iowa us-central1; retrieved 25 October 2026.