Scarcity, allocation, and total reliance

A completed synthetic confirmation under a shared rate budget

Evidence audit and working manuscript v0.2 • 6 November 2026 • Prepared for Ben Amuwo. AI-assisted, author-commissioned audit; no independent human peer review or experimental replication is claimed.

Abstract

The RC4 confirmation tested whether internally generated source-monitoring channels in the declared U2, eta=0, Arm S, lead k=4 fixture exhibit registered scarcity, displacement, total-reliance and allocation-prediction relations. A fixed panel paired 501 training seeds across six environment-by-capacity conditions. All 3,006 conditions reached 28,000 training steps. The primary confirmation population retained every seed; a component failing its registered validity prerequisites contributed zero by definition. [S01-S05]

A retrospective reconstruction from the final assembled data reproduces confirmation of T1, T2, U and V and the recorded overall PASS. Composite normalized scarcity losses were 0.423863 in base and 0.354734 in shift; source 2 gained 0.497778 of total rate share. All six primary replacement cells exceeded the 0.02 margin, and all eleven allocation comparisons favoured reference-v2 under the registered Holm family. These are local findings under the frozen procedure, not architecture-general laws or evidence of strict metarepresentation or ecology.

The frozen analyzer replay produced a byte-identical adjudication. A separately written, source-aware arithmetic audit matched 481 counts, summaries, seed-invalidity lists, inferential values and labels within stated numerical tolerances. The execution record also contains an unplanned stop after an environment-distribution verification refusal; the cause and remediation remain unverified. The statistical result is reproducible from the assembled data, while execution-history assurance remains qualified.

Record	Current assessment
RC4 statistical disposition	PASS reproduced from final assembled data
Numerical and local integrity checks	Supported within the audit boundaries
Historical interruption	Disclosed; recovery provenance unresolved
Manuscript / public release	Public working draft; no peer-review claim

A reading guide: three useful monitors

Imagine three instruments sharing a narrow communication link. Each knows something private. A decision-maker sees public context and their messages. The question is how a limited shared information budget changes what is transmitted and used. This is an analogy for a finite categorical experiment, not a claim about biological attention or energy.

Term	Operational meaning	Distinction to retain
C	Shared capacity target in bits	Final gates allow registered slack.
r_v	Expected KL rate for source v	Includes learned-prior mismatch.
a_v	r_v divided by total realized rate	A share, not a return or raw bit count.
J	Expected additive task return	Defined by fixture reward weights.
H	Full-information minus public-only return	Fixed headroom for normalization.
Seed	One paired training initialization	Six correlated conditions per seed.
Gate	Registered measurement prerequisite	Completion alone is insufficient.
Composite	Raw value when valid, zero otherwise	All 501 positions remain in the mean.

Four questions, four bounded answers

T1 asks whether a smaller budget costs return. T2 asks whether a more valuable source gains rate share. U asks whether replacing a message independently from its context-conditioned marginal costs return. V asks whether a frozen allocation prediction is closer to observed shares than specified comparators. U is primary at high capacity; the low-capacity diagnostic must not be substituted for it.

How to read the visual companion

The local website companion separates archived measurements, arithmetic on an existing selected seed, and a dashed teaching panel with invented values. The seed selector does not generate data. Allocation bars show means of per-seed shares. The prediction triangle shows share coordinates; its screen distances are not the L1 error used for inference. The reported decision is over the registered population, not one selected seed.

Version 0.2 adds this reading guide and accompanies the E1 visual research record. The scientific disposition, effect estimates, gates and unresolved interruption are unchanged from the 6 November audit. It remains a working paper. No external archival deposit or independent human validation is implied.

The construct and the question

The question is local: when several useful internal-state monitoring channels share a finite communication budget, do their returns and allocations change predictably with capacity and source value? The experimental unit for inference is the training seed, not a timestep, a sampled task trial or one of the six correlated conditions. [S02-S04]

Finite synthetic fixture

The inspected implementation uses three private source alphabets of sizes 4, 2 and 3. Public context I consists of a binary context X and public observations Z1, Z2 and Z3. In the declared eta=0 model the source posteriors factor conditional on I. Public reliability parameters are rho=(0, 0.25, 0.5), with P(Cv=c | Zv=z)=(1-rho_v)/n_v + rho_v 1[c=z]. Contexts are enumerated uniformly. [S09]

Encoder v has source-local access q_v(Sv | Cv,I), with independent private randomness. The consumer may read I and all messages. Reward is additive: the sum over sources of the context-dependent weight for a correct coordinate action. The adopted source-local restriction matters: conclusions about this class do not automatically apply to centralized encoders that read every private source. [S09,S10]

Environment	Weights at X=0	Weights at X=1	Full / public return
base	1, 1.8, 1.1	2, 0.9, 1.6	4.2 / 2.11875
shift	1, 7.2, 1.1	2, 3.6, 1.6	8.25 / 4.65
alpha	1, 9, 1.1	2, 4.5, 1.6	9.6 / 5.49375
beta	2.5, 1.8, 2.2	5, 0.9, 3.2	7.8 / 3.58125

Shared rate accounting

```
r_v = E[ KL_bits( q_v(. | C_v,I) || p_v(. | I) ) ]
R = r_1 + r_2 + r_3; a_v = r_v / R when R > 0
```

The learned prior p_v is distinct from the induced message marginal q_v(Sv|I). The KL rate includes mutual information and prior-mismatch cost; the total is shared across channels. Training uses a common dual multiplier and a scheduled capacity target. Final validity allows R <= C+0.02, rather than requiring exact equality with nominal C. [S02,S09]

This section documents the inspected implementation and frozen definitions. It is not a new clean-room specification and must not be transferred into unfinished R1 design without a separately governed handoff.

Design and training procedure

The frozen registry contains seeds 9201-9701, paired in a fixed seed-major order across base C=4.5, base C=0.8, shift C=4.5, shift C=0.8, alpha C=0.8 and beta C=0.8. No seed replacement or adaptive sample extension is part of the registered design. [S02,S03]

Item	Registered implementation
Training length / batch	28,000 steps / 256
Warm-up	0-3,999: capacity target 4.5; balanced reward weights
Capacity ramp	4,000-15,999: linear transition to target capacity
Hold	16,000-19,999: target capacity
EMA phase	20,000-27,999: decay 0.9997; learning-rate decay
Learning-rate multiplier	(1 + max(0, step-20000)/1000)^(-2)
Optimizer / clipping	Adam; global gradient norm clipped at 10
Base learning rates	0.035, 0.004, 0.025, 0.035, 0.004 by parameter group
Dual update	max(0, dual + 0.005*f*(sampled rate-target))
Checkpoint / exact evaluation	Every 2,000 / every 500 steps; final scoring uses EMA

The active trainer is MB1-PCR4-v1: hard categorical messages, a Rao-Blackwellized score-function encoder reward gradient, and exact categorical KL gradients conditional on the sampled context and private state. It is not the older Gumbel-Softmax update simply because that older module remains in the import closure. The consumer gradient uses sampled hard-message utility. Shared and full-context residual parameter tables are optimized jointly. [S09]

Why 501?

Sample size was selected for a prospectively stated analyzability objective. For the V environments, a one-sided exact upper failure bound from 0 failures in 8 pilot seeds was 0.3123439781. With at least 324 valid seeds required per V pair, the dependence-free joint lower assurance was 0.9483276870 at n=500 and 0.9554441273 at n=501. The prior reviewer calculation reproduces today. The 2.0-fold standard-deviation transport multiplier is a planning assumption, not an estimated population law or a guarantee of statistical power. [S07,S11]

Estimands, gates, and decisions

```
H_base = 4.2 - 2.11875 = 2.08125
H_shift = 8.25 - 4.65 = 3.6
T1_e(s) = [J_e,4.5(s) - J_e,0.8(s)] / H_e
T2(s) = a_shift,0.8,2(s) - a_base,0.8,2(s)
U_e,v(s) = [J_e,4.5(s) - J_replaced,e,v(s)] / H_e
V_j(s) = ||a_s - b_j||_1 - ||a_s - b_reference-v2||_1
```

For U, replacement draws Sv independently from the trained induced marginal q(Sv|I), holding world state, I and the jointly generated other messages fixed, then recomputes the trained consumer return. Primary U is evaluated at C=4.5. Replacement at C=0.8 is a descriptive diagnostic, not an alternative primary test. [S02,S03,S09]

Condition gates are recomputed from assembled fields: completion at 28,000 steps; finite values; positive total rate and agreement of component rates with the stored total to 1e-9; budget within 0.02; and a strict final-nine evaluation range below 0.002*H. High-capacity base/shift additionally require the reference upper value to match the frozen full return and full-J <= 0.01*H. U adds a source-specific active-information gap <=0.01 bits and replacement-original agreement within 5e-12. [S02,S05]

T1 requires both paired conditions valid; T2 requires both low-capacity training conditions; V requires the relevant alpha/beta low-capacity condition. Invalid components contribute exactly zero to all-seed confirmation. Valid-only disconfirmation/equivalence requires at least 20 observations for T1/T2/U and 324 for V; otherwise p=1 retains the frozen family size.

At alpha=0.05, confirmation uses one-sample Student t inference: both T1 environments exceed 0.05, T2 exceeds 0.05, all six U cells exceed 0.02, and all eleven V advantages exceed 0.05 with Holm correction. T1/U conjunctions use intersection-union logic. Valid-only disconfirmation uses Holm across the two T1 or six U cells, T2 an upper bound below zero, and V separate eleven-pair comparator-superiority and equivalence families. Simultaneous confirmation and disconfirmation yields INCONCLUSIVE.

All four families CONFIRMED -> PASS; any DISCONFIRMED -> FAIL; otherwise HOLD, including structural refusal. Student t inference carries its sampling/distributional assumptions. The archived lower_95 and upper_95 are individual one-sided 95% bounds; together they are a central 90% interval, not a two-sided 95% interval.

Scarcity and value-sensitive displacement

Component	Valid n	Composite mean	Lower 95%*	Margin
T1 base	381/501	0.423863	0.406330	0.05
T1 shift	429/501	0.354734	0.344012	0.05
T2	501/501	0.497778	0.497231	0.05

*One-sided, unadjusted 95% t lower bound. All composite estimates retain n=501. Values rounded to six decimals here; full-precision values and tolerances are in the audit files.

The two normalized high-to-low-capacity return losses exceed the 0.05 margin under the registered composite analysis. T1 retained 381 valid paired seeds in base and 429 in shift. The other 120 and 72 seeds respectively contributed zero, rather than being dropped. These estimates describe the registered composite outcome, not the hypothetical effect among universally well-optimized runs.

All 501 low-capacity base/shift pairs met their condition gates. Source 2 gained approximately 49.78 percentage points of total rate share. The mean changes in sources 1 and 3 were -0.404795 and -0.092982 respectively. Their sum offsets the source-2 increase; the largest per-seed conservation residual was 2.36e-16. This is displacement within a normalized shared allocation. [S04,S05,A02]

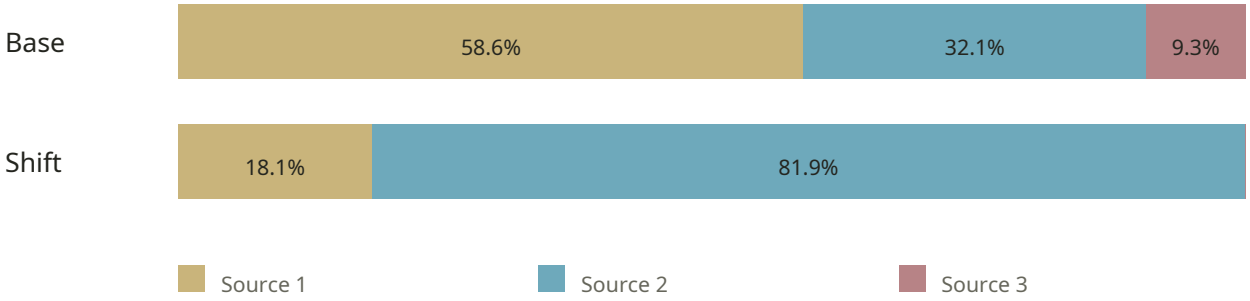


Figure 1. Mean rate shares over all 501 paired seeds at C=0.8. Bars are descriptive averages; the registered T2 inference uses each seed's paired source-2 difference. This figure does not introduce a new hypothesis test.

Completion is not validity

Every scheduled condition completed. Nonetheless, base C=4.5 had 118 nonbinding-gate failures and 2 plateau failures; shift C=4.5 had 70 nonbinding-gate failures and 2 plateau failures. Every low-capacity condition was valid. The report therefore distinguishes zero terminal execution failures from nonzero scientific validity failures.

Total reliance: six primary cells

Primary U cell	Valid n	Composite mean	Lower 95%*	Margin
base / source 1	381/501	0.411070	0.394068	0.02
base / source 2	381/501	0.231227	0.221663	0.02
base / source 3	381/501	0.246642	0.236441	0.02
shift / source 1	402/501	0.250748	0.241578	0.02
shift / source 2	429/501	0.602077	0.583899	0.02
shift / source 3	407/501	0.152320	0.146926	0.02

*One-sided 95% t lower bound; the six-cell confirmation uses the registered intersection-union rule. Every composite retains 501 seed positions. Primary intervention capacity is C=4.5.

All six cells exceed the registered 0.02 margin. In addition to the high-capacity condition gates, source-specific active-channel prerequisites reduce shift/source 1 to 402 valid seeds and shift/source 3 to 407. Shift/source 2 retains 429. Validity is consequently component-specific: one source's invalidity does not erase another source's admissible measurement. [S04,S05,A02]

The audit recomputed the primary losses from the assembled original and replaced returns. It also checked internal diagnostic identities: target-specific returns sum exactly to the stored totals; original replacement-path returns agree with condition returns to 2.67e-15; and unique-versus-total replaced returns differ by at most 3.56e-15. These are consistency checks on supplied readouts, not new enumeration from checkpoint parameters.

A boundary the paper must preserve

The descriptive low-capacity shift/source-3 replacement loss averages approximately 0.00001746, despite confirmation of the primary high-capacity source-3 claim. That diagnostic cannot change the registered disposition, but it prevents a broader claim that every source remains strongly relied upon at the binding budget. Allocation toward the high-value source can coincide with very little low-capacity reliance on another source. [S05]

Total-reliance replacement establishes a cost under the declared intervention and trained consumer. It does not identify strict metarepresentational content, prove an unconditional mechanism effect across failed optimizations, or establish general success beyond this fixture. The source-local independent model also makes the unique-reliance diagnostic coincide with total reliance; that identity must not be advertised as a discovery about correlated systems.

Allocation prediction: eleven comparisons

Env.	Comparator	Mean advantage	Lower 95%*	Holm p
alpha	Uniform	1.059535	1.057175	<1e-300†
alpha	Conditional entropy	1.272364	1.270004	<1e-300†
alpha	Inferability only	1.030581	1.028221	<1e-300†
alpha	Isolated causal value	0.493325	0.490965	<1e-300†
alpha	Proportional one-step value	0.414438	0.412078	<1e-300†
beta	Uniform	1.003107	1.000265	<1e-300†
beta	Conditional entropy	0.718766	0.715924	<1e-300†
beta	Inferability only	0.940944	0.938102	<1e-300†
beta	Isolated causal value	0.336441	0.333599	<1e-300†
beta	Proportional one-step value	0.176319	0.173540	3.148e-274
beta	Training-environment proxy	0.491616	0.489604	<1e-300†

All eleven pairs have 501/501 valid seeds and margin 0.05. *Displayed t bounds are one-sided and unadjusted; decisions use the Holm-adjusted p-values. †The archived calculation returns floating-point zero; the display conservatively uses <1e-300, not an assertion of probability exactly zero.

Every pair is classified v2_superior. The smallest mean advantage is 0.176319 for beta against proportional one-step value, with Holm-adjusted p approximately 3.15e-274. This is evidence for the frozen reference against these eligible comparisons, not every conceivable allocation model. [S02,S04,S05]

The contract's alpha and beta reference-share vectors match the adopted reference-v2 artifact exactly, and the four adopted artifact hashes match their freeze record. A separate historical AI technical review reports proof/certificate and high-precision checks. This audit verifies that provenance link and the downstream comparisons; it does not freshly regenerate the interval certificates. [S10,A04]

Reference-v2 was adopted only within its stated k=4, C=0.8 scope. The historical midpoint-return discrepancy remains disclosed and is not rewritten. The scientific effect margin, comparator-eligibility tolerance and certificate uncertainty are different quantities and must remain separate.

What the audit actually verifies

Check	Observed result	Boundary
Canonical/package/authority hashes	95/95 comparisons match; 42 package entries; no extras	Identity, not external authentication
Final-data structure	501 ordered seeds; six conditions each; schema valid	Supplied assembled data
Ledger	177,356 entries; no chain/index breaks	Self-recorded chain
Execution schedule	42,084 starts and checkpoints; all 3,006 complete in order	Recorded events
Checkpoint references	All final references match ledgered plaintext digests	No checkpoint decryption
Frozen environment verifier	8/8 checks true; empty distribution diff today	Current state only
Frozen analyzer replay	Exit 0; adjudication bytes identical	Same code and runtime
New numerical reconstruction	481 comparisons; no discrepancies at stated tolerances	Source-aware; shared SciPy
Synthetic analyzer tests	18/18 pass; no training or runner initialization	Software fixtures only
Backup payload	12,099/12,099 listed files match; 12,024 ciphertext links match	Protected directory excluded

The final data names the ledger head immediately before assembly. The ledger then appends one data_assembled event naming that data hash. The two head digests therefore differ as designed; the pre-assembly link and committed output hash both match. Treating this expected difference as corruption would be an audit error. [S04,S06,A01]

The current environment check validates the locked interpreter, platform, distribution metadata, module identities, torch build, vendored scientific sources and determinism exports. Its 197-distribution lock hashes RECORD metadata where available; jsonpointer and zstandard lack such a record. It is not a fresh hash of every installed library payload, nor proof of the machine's historical state at every step.

The prior CF-4 reviewer identifies itself as a fresh Kimi Work AI agent and reports 38 development/test checks. Today's audit reruns the 18 synthetic analysis tests and the review's sample-size arithmetic, but no training, runner, decryption or integration tests. The historical 38/38 statement must not be relabelled as a new 38-test replication. [S07,A04]

Deviations, qualifications, and open evidence

Unplanned environment refusal and restart

A retained control record reports five consecutive driver errors caused by GuardRefusal: environment verification failed for distributions, at 20:43:12 on 22 October 2026 in Vancouver (03:43:12 UTC, 23 October). A subsequent driver_started record is timestamped 20:54:11 local. The author states on 6 November that this was an unplanned stop and that they do not know what happened. [S08,A05]

The record supports a fail-closed refusal and a later restart record. It does not identify the changed distribution, cause, remediation, operator actions, contemporaneous verifier report or recovery authorization. The final ledger has the expected complete condition/chunk schedule and no recorded runner exception/replay event, but a driver-level failure need not appear as a scientific runner event. Do not claim uninterrupted execution, no environment drift, or fully verified recovery.

Independence and chronology

The pre-execution CF-4 review and earlier reference review are AI reviews; no unaffiliated human peer review is established by the inspected record. Commitment, access and execution timestamps are self-recorded. Hash linkage is strong local integrity evidence but does not independently establish when a file first existed or prove absence of interim access. The publication should use internally frozen/prospectively specified language with those limitations, rather than implying independently timestamped public preregistration.

Preservation and publication packaging

The recovery backup README explicitly says protected decryption keys are included. Its protected directory was not opened, copied or hashed. The audited 12,099 manifest-listed files exclude that protected area. The backup is a same-host recovery copy; the inspected evidence does not establish independent off-device disaster recovery. It must not be uploaded wholesale as a public research artifact.

The local review bundle instead contains selected final data, frozen package, authority records, ledger and audit receipts. It excludes private custody and sealed payloads. The retrospective audit did not alter the original scientific records. The author subsequently authorized release of this public working draft and companion on sidus.quest; no independent validation or journal acceptance is implied. The R1 planning snapshots remain unchanged; this context is outcome-exposed and no longer eligible as an unexposed clean-room role.

Interpretation and further validation

The most defensible present statement is that the recorded RC4 all-seed confirmation result is computationally reproducible from its assembled data under its frozen analysis rule. It supports local scarcity, value-sensitive displacement, registered high-capacity replacement reliance, and reference-v2 allocation advantage in the declared synthetic fixture. It does not establish an independent replication or certify every stage of original data generation.

The observed optimization/gate failures are part of the primary outcome definition, not details to hide. Valid-population analyses are conditional on successful optimization and scientific gates; they do not identify an unconditional mechanism effect. No secondary diagnostic, changed threshold, pooled analysis or future R1 result may retroactively rescue or broaden the original label.

Before calling this a finalized paper

Work item	Evidence or decision still required
Execution-history closure	Preserve the unexplained refusal/restart; locate recovery evidence if possible, otherwise disclose the unresolved limitation without claiming full recovery verification.
Methods and lineage review	Check the complete source-local construct, training estimator and reference certificate argument; distinguish code inspection from proof and execution-level reproduction.
Independent critique	Seek bounded external statistical/methods review; no institutional endorsement is invented or required to describe the author's result.
Author sign-off	Approve authorship, contribution/AI disclosure, prior development chronology, deviations and precise scope.
Public release	Approve rights/licenses, citation details and a private-key-free versioned archive; publish only the reviewed release.
R1 separation	Keep detailed RC4 outcomes outside any unfinished clean-room specification/implementation context.

Report status

This is a substantive working report and evidence audit, not a submission-ready final manuscript or a new scientific adjudication. The historical PASS is preserved and reproduced. The interpretation remains qualified by the open execution-history issue and the stated scope of verification. No new training is required to describe these findings honestly; additional validation would be labelled by what it actually establishes.

Presentation follows the linked E0-A/E0-B records in separating the question, operational definitions, results, provenance, and scope limits. Their scientific findings were not revalidated in this RC4 audit; the local site copies were used as presentation references when browser extraction was incomplete.

SUPPLEMENT / TRACEABLE RECORD

Evidence index and reproducibility

ID	Evidence used	Full local audit-bundle location
S01	Canonical claim freeze and freeze manifest	evidence/canonical/
S02	Frozen machine contract and schema	evidence/rc4_package/contract/
S03	Protocol, build decisions and schedule	evidence/rc4_package/PROTOCOL_RC4.md
S04	Final assembled confirmation data	evidence/final/CONFIRMATION_DATA_RC4.json
S05	Recorded adjudication and completion	evidence/final/
S06	Ledger, public status and run lock	evidence/run/
S07	CF-4 review, results, reproduction log	evidence/governance/21_CF4_.../
S08	Driver exception/start/completion records	evidence/control/
S09	Frozen trainer, evaluator and readout	evidence/rc4_package/vendor/ and rc4/
S10	Reference-v2 adoption and AI review	evidence/reference_v2/
S11	Sample-size planning records	evidence/planning/
A01	Fresh local integrity/ledger audit	INTEGRITY_AUDIT.json
A02	New numerical reconstruction and tables	NUMERICAL_AUDIT.json; PRIMARY_RESULTS.csv
A03	Frozen analyzer and environment replay	FROZEN_ANALYZER_REPLAY_RECEIPT.json
A04	Tests, sample size, reference identity	VERIFICATION_RECHECKS.json and logs
A05	Open incident and author account	OPERATIONAL_INCIDENT_REGISTER.json
A06	Backup payload hashes and limitations	BACKUP_PAYLOAD_AUDIT.json

The full local audit bundle (retained separately from the selected website evidence) contains these records. Its SOURCE_INDEX.json contains exact copied-source paths, byte sizes and SHA-256 identities. DELIVERABLES_MANIFEST.json separately binds the local report and audit files. The archived adjudication and replay share SHA-256 aea55869cb72929d0622c55626740f2033506e3e7e0fd23db5fbd3d01ca5386d. A hash is an identity claim, not an external timestamp or scientific endorsement.

```
python reproduce_numeric_audit.py \
--package-root evidence/rc4_package \
--final-root evidence/final \
--output-dir reproduced
```

The new numerical audit requires Python, SciPy and jsonschema; it does not import RC4 code, train, decrypt or assemble outcomes. It is source-aware, retrospective and shares SciPy with the original analysis. Use the included README for recorded versions, tolerance policy and the distinction from an exact-environment frozen-analyzer replay. Readme commands do not authorize another confirmation run.