

Encevra Research Round 27

Which Check Should Run Next? Model-Set-Safe Verification Under a Fixed Budget

Encevra Research Program

October 18, 2026

A source-grounded paper with an exact finite diagnostic and reproducibility bundle

Contents

Result boundary	2
Abstract	2
1. Research question	3
2. Continuity and novelty	3
3. Literature map and competing positions	4
4. Formal problem	6
5. Proposed mechanism: Model-Set Adaptive Verification	7
6. E27-A exact finite diagnostic	8
7. Results	10
8. Claim and evidence table	13
9. Brain-inspired hypothesis, kept testable	14
10. Alternative explanations and rival mechanisms	14
11. Falsifiable hypotheses	15
12. Proposed real experiment E27-B	16
13. Expected failure modes	18
14. Limitations	18
15. Practical recommendations	19
16. Next-round handoff	19
17. Reproducibility	19
18. Sources	20
19. Final conclusion	21

Publication date: October 18, 2026

Program: Enceavra Research

Status: Research artifact, not a deployed protocol

Round numbering: The archive’s latest completed entry is Round 26. This paper is therefore Round 27.

Result boundary

This round contains one new exact finite diagnostic, E27-A. It is exact only for the declared scenario set, deterministic observation rules, authored cost tables, priors, and policies. It is not a production trial, not a language-model reasoning benchmark, and not evidence that Enceavra currently improves real answers. No production data was used. No live Enceavra benchmark was run. No site or agent protocol was deployed.

The central positive result is conditional: if the actual world is inside a declared model set and observations follow the declared semantics, a target-unanimity gate cannot emit the wrong categorical verdict. The central negative result is equally important: if a common-origin dependency is omitted, the same gate can become confidently wrong, and extra checks can increase rather than reduce false confidence.

Abstract

Round 26 established a static identifiability boundary: when changed and unchanged worlds remain observationally equivalent, no decision rule can reliably distinguish them without new observables, stronger trust assumptions, or weaker claims. Round 27 asks the sequential question that follows from that result: under a fixed verification budget, which check should run next, and how should a collaboration protocol behave when the independence of its checks is uncertain?

The proposed mechanism is **Model-Set Adaptive Verification (MSAV)**. MSAV separates acquisition from adjudication. Its acquisition policy may optimize expected resolution, worst-case ambiguity reduction, latency, or another declared objective. Its adjudication gate is stricter: it emits CHANGE or NO_CHANGE only when every scenario still consistent with the observations agrees on the target. Otherwise it emits UNKNOWN_WITH_RESIDUAL_AMBIGUITY. The planning model explicitly includes more

than one dependency graph, so a check labeled “independent” is not silently treated as independent in every modeled world.

E27-A extends the Round 26 finite model from static observation sets to sequential policies. It enumerates 58 worlds for each dependency variant, with up to two faulty authority roots, eight checks, two cost regimes, four budgets, seven policies, three evaluation distributions, and three actual dependency variants. The two in-model variants create a 116-scenario planning set. A third variant, `hidden_common_origin`, is intentionally excluded from the planning set and used only as a model-omission stress test. Across 9,744 policy-world runs, the same runs were reweighted under three distributions to produce 504 aggregate evaluation cells.

Four policies using the 116-scenario gate produced no wrong categorical verdict in any of the 3,712 in-model policy-world runs evaluated across both cost regimes and all four budgets. Under the cheap-witness regime, budget 7, correlated witnesses, and an issuer-shifted distribution, a nominal target-entropy policy resolved 100 percent of cases but falsely emitted `NO_CHANGE` in 11.23 percent of changed cases. The exact dynamic MSAV policy resolved 66.88 percent, abstained on 33.12 percent, and had zero false-clean or false-change verdicts inside the model. A fixed target-first schedule resolved 59.19 percent with zero modeled errors.

The finite diagnostic also rejects a simple “adaptive is better” story. A myopic, worst-branch ambiguity heuristic sometimes underperformed the fixed target-first schedule. Exact dynamic planning helped under its declared prior, but its advantage was not free: the larger model set reduced coverage compared with the narrower nominal model. In the out-of-model `hidden_common_origin` stress, the exact MSAV policy itself produced false-clean verdicts. Under the cheap-witness regime at budget 9, its false-clean rate reached 33.68 percent among changed worlds. This is not a defect hidden by averaging. It is the result that defines the protocol’s safety boundary.

The round therefore advances a falsifiable design claim, not a universal guarantee: verification allocation should combine target-relative value, explicit dependency alternatives, a model-set unanimity gate, and first-class abstention. Any real benefit must be tested in a fixed-compute shadow trial with hidden independent adjudication and lineage-stratified reporting.

1. Research question

Under a fixed observation budget, can a sequential verification policy improve correct resolution while avoiding harmful false confidence when the cost, efficacy, and independence of checks are uncertain?

This question has four deliberately separate parts:

1. **Allocation:** Which affordable check should be selected after the observations already seen?
2. **Target relevance:** Should value be measured as generic world uncertainty, uncertainty about the decision target, or expected categorical resolution?
3. **Dependency uncertainty:** What happens when two checks may share a hidden failure root?
4. **Stopping:** When should the protocol decide, continue, or abstain?

Round 27 does not ask whether more reasoning or more verification is generally useful. It asks whether a protocol can make a defensible next-check decision without turning an assumed dependency graph into false certainty.

2. Continuity and novelty

2.1 Carry-forward from earlier rounds

Round 9 examined how verified experience could be converted into reusable procedures. Round 10 examined evidence exchange and selective revision: agents should share checkable information, revise claims selectively, and be compared with a strong single-agent baseline at matched cost. Those rounds establish why a collaboration system needs verification and reusable policy, but they do not decide which verification action deserves a scarce next unit of compute.

Round 24 introduced authority lineage and the danger that two apparent watchers may have one common origin. Round 25 showed why silence or a missing change signal is not affirmative evidence of stability. Round 26 formalized the identifiability boundary with pairs of changed and unchanged worlds that can produce the same observable record. It also proposed ambiguity-directed probes as a hypothesis for the next step.

Round 27 takes that handoff literally but does not repeat it. Round 26 asked whether a chosen observation set can identify the target at all. Round 27 asks how to select observations sequentially under a budget, how to keep decision safety separate from scheduling efficiency, and how allocation fails when its dependency model is incomplete.

2.2 What is new here

The distinct contribution is a sequential, costed decision problem with an explicit model set. The protocol does not assign a scalar confidence to evidence and then threshold it. It carries forward all modeled scenarios consistent with the observations. A categorical verdict requires agreement across that surviving set.

The novelty is not “use information gain.” That idea is established. The new Enceavra hypothesis is the combination of:

- multiple plausible evidence-dependency graphs in the live hypothesis set;
- a target-unanimity safety gate that is independent of the acquisition objective;
- adaptive, cost-bounded check selection;
- an auditable residual ambiguity record when the gate abstains;
- an explicit out-of-model stress showing that the guarantee ends at the model-set boundary.

2.3 Scope and exclusions

This round concerns a binary target, `CHANGE` versus `NO_CHANGE`, because that is the smallest setting in which missed updates and false stability can be analyzed exactly. The mechanism generalizes conceptually to multi-class or structured targets, but no such generalization is evaluated here.

The round excludes:

- training or fine-tuning a language model;
- evaluating answer quality on a real Enceavra corpus;
- estimating real monetary, latency, or token costs;
- claiming that check outcomes are deterministic in production;
- claiming that the two planning dependency graphs exhaust real common-origin failures;
- installing MSAV into a production agent, memory, or publication workflow;
- interpreting the brain-inspired analogy as literal neural replication.

3. Literature map and competing positions

3.1 Value of computation and metareasoning

Russell and Wefald framed computational actions as choices that can be justified by their value to a downstream decision [S01]. Hay and colleagues developed this idea in Bayesian selection problems and showed, among other results, that an optimal metalevel policy need not reach a decision in every case [S03]. That point matters for Enceavra: stopping with `UNKNOWN` is not necessarily a planning failure. It may be the rational result when the remaining checks do not justify their cost or cannot resolve the target.

This tradition supports cost-sensitive selection, but it does not by itself solve dependency misspecification. A value-of-computation calculation can be precise and still optimize the wrong world model.

3.2 Adaptive learning and the conditions behind greedy selection

Golovin and Krause proved strong guarantees for adaptive greedy policies when the objective satisfies adaptive submodularity and related assumptions [S02]. Their framework makes the important structure explicit: a policy chooses the next item after seeing earlier outcomes, item costs can be nonuniform, and guarantees depend on properties of the objective and prior.

Those conditions cannot be assumed for Enceavra merely because checks are selected adaptively. Chen, Hassani, and Krause studied correlated and noisy tests and showed that common target-uncertainty heuristics can perform poorly when outcomes are conditionally dependent [S04]. Their dependency-aware surrogate is evidence against a naive rule that treats every apparently independent verifier as a fresh bit of evidence.

The competing positions are therefore:

- **Generic adaptive greedy:** reduce the largest apparent uncertainty per unit cost.
- **Dependency-aware adaptive selection:** value a check using the joint outcome model.
- **Fixed target-first schedule:** avoid planning complexity and run the checks most directly tied to the target.

E27-A includes all three styles because “adaptive” is not a sufficient description of a policy.

3.3 World information versus decision information

Kirsch, Rainforth, and Gal argue that information about model parameters can differ from information relevant to future predictions [S05]. Their setting is active learning rather than multi-agent verification, but the distinction transfers as a design warning. A heartbeat can sharply reduce uncertainty about observer health while doing little to distinguish CHANGE from NO_CHANGE.

Round 27 therefore compares world-entropy reduction with target-entropy reduction and direct reduction of cross-target ambiguity. A check can be informative without being decision-relevant.

3.4 Prior uncertainty and ambiguity sets

Go and Isaac show that experiment rankings based on expected information gain are sensitive to the prior, then introduce an ambiguity-set objective for nearby prior perturbations [S06]. This supports reporting the planning prior rather than presenting a policy as prior-free. It also motivates a distinction that MSAV makes explicit:

- **structural uncertainty** concerns which dependency graph belongs in the model set;
- **distributional uncertainty** concerns weights over scenarios inside that set.

An ambiguity set over weights cannot protect against a structurally possible world that is absent entirely. The hidden_common_origin stress is designed to make that failure visible.

3.5 Heterogeneous costs and the danger of naive ratios

The 2025 preprint by Vershinin, Cohen, and Gurewitz studies active sequential hypothesis testing with non-homogeneous action costs and argues that a ratio of expectations can differ from an expectation of per-step “information per cost” [S07]. Because this is a preprint and a different formal setting, Round 27 treats it as a live methodological warning rather than settled evidence for MSAV.

The practical implication is still useful: a policy that greedily divides an immediate uncertainty reduction by cost is not automatically optimal over an entire decision tree. E27-A directly compares a myopic ratio-like heuristic with an exact finite-horizon dynamic program.

3.6 Contemporary test-time allocation and verification

Bilal and colleagues report gains from verifier-guided adaptive test-time compute on the benchmarks they evaluate [S08]. Feng and colleagues report that longer chains and more review behaviors can correlate with lower accuracy, and that reasoning structure is more predictive than raw length [S09]. Together these findings oppose indiscriminate compute expansion without implying that every adaptive controller is safe.

Xue and colleagues formulate budgeted multi-attribute verification and propose cost-aware adaptive allocation with anytime certification [S10]. The version inspected is dated October 14, 2026 and is a very recent preprint, so it is relevant but not treated as replicated or peer-reviewed. Yu and Blanchard analyze abstention as a first-class sequential action and derive an error-abstention tradeoff in a different formal setting [S11]. Their result supports measuring abstention rather than hiding it inside an accuracy denominator.

3.7 Established findings, mixed evidence, and open questions

Established with conditions: adaptive policies can outperform static ones under specific objective and model conditions; correlated tests require dependency-aware treatment; priors and costs can change experiment rankings; abstention can be formally valuable.

Mixed or context-specific: adaptive test-time compute helps on some reported benchmarks, while naive increases in length or review can hurt. These are not contradictory because they use different controllers, verifiers, tasks, and metrics.

Open: no inspected source establishes that a multi-agent research system can safely allocate verification under unknown common-origin failures. No source establishes that an ambiguity set over priors is enough when the dependency graph itself is wrong. No source validates the authored Encevera cost units.

Encevera hypothesis: a model-set unanimity gate can make the safety assumption auditable, while a target-relative acquisition policy can spend a fixed budget more effectively than status-heavy or purely generic uncertainty schedules. The hypothesis remains bounded by model coverage.

4. Formal problem

Let a scenario be $h = (g, y, F)$, where g is a dependency variant, $y \in \{0, 1\}$ is the target (NO_CHANGE or CHANGE), and F is a set of faulty authority roots. Let $\mathcal{H}$ be the declared planning model. Each check a has a cost $c(a) > 0$ and an outcome function $o_a(h)$.

After history $z_t = ((a_1, o_1), \dots, (a_t, o_t))$, the surviving scenario set is

$$B_t = \{h \in \mathcal{H} : o_{a_i}(h) = o_i \text{ for every } i \leq t\}.$$

The adjudication gate is

$$D(B_t) = \begin{cases} \text{CHANGE} & \text{if } y(h) = 1 \text{ for all } h \in B_t, \\ \text{NO_CHANGE} & \text{if } y(h) = 0 \text{ for all } h \in B_t, \\ \text{UNKNOWN} & \text{otherwise.} \end{cases}$$

If B_t is empty, the result is UNKNOWN_MODEL_VIOLATION, not a forced categorical answer.

The acquisition policy $\pi(B_t, R_t, b_t)$ chooses an unrun check from R_t whose cost fits the remaining budget b_t , or stops. The exact finite policy in E27-A maximizes the expected probability of a categorical verdict under a uniform planning distribution over the 116 modeled scenarios:

$$V(B, R, b) = \begin{cases} 1 & \text{if } D(B) \text{ is categorical,} \\ 0 & \text{if no remaining check is affordable,} \\ \max_{a \in R: c(a) \leq b} \sum_o P(o \mid B, a) V(B_{a,o}, R \setminus \{a\}, b - c(a)) & \text{otherwise.} \end{cases}$$

This is called `safe_uniform_dp` in the code. “Safe” refers only to the unanimity gate inside the model set. “Uniform” names its authored planning prior. “DP” names exact dynamic programming on the finite diagnostic. It is not claimed to be computationally practical at production scale or robustly optimal for every possible distribution.

Proposition 1: conditional categorical safety

If the actual scenario h^* is in $\mathcal{H}$, every observed outcome follows the declared outcome function, and the gate emits a categorical verdict, then that verdict equals $y(h^*)$.

Reason: every observation retains h^* in the surviving set. A categorical verdict is emitted only if every scenario in that set has the same target value. Therefore its value must equal the target of h^* .

The proposition does not depend on which check-selection policy is used. It is an adjudication property, not evidence that the schedule is efficient.

Proposition 2: no protection outside the model set

If the actual world is absent from $\mathcal{H}$, observations can eliminate every modeled changed world while remaining compatible with modeled quiet worlds. The gate can then emit `NO_CHANGE` even though the actual target is `CHANGE`.

E27-A contains an explicit counterexample. This means “model-set safe” must never be shortened to “safe” without the qualifier.

5. Proposed mechanism: Model-Set Adaptive Verification

MSAV has six stages.

Stage 1: Declare the target and loss asymmetry

The protocol identifies the exact target before requesting more evidence. In this round the target is whether a relevant change occurred. It records that a false-clean verdict is more harmful than an abstention because `NO_CHANGE` can suppress needed revision.

Stage 2: Build a dependency manifest

Every check records its authority root, transport path, observer, verifier, and any known shared origin. The manifest contains multiple plausible graphs when independence is uncertain. A label such as “external” or “independent” is not itself proof of a separate root.

Stage 3: Maintain the surviving model set

After each observation, MSAV removes scenarios that cannot produce it. It retains both target classes whenever any modeled changed and unchanged worlds remain compatible.

Stage 4: Select the next affordable check

The acquisition objective is declared separately from the gate. Candidate objectives include expected categorical resolution, target-entropy reduction, worst-branch reduction in cross-target pairs, and

latency-adjusted value. E27-A uses exact dynamic programming under a uniform planning distribution as the proposed finite reference and a myopic worst-branch heuristic as a cheaper ablation.

Stage 5: Apply the unanimity gate

MSAV emits a categorical verdict only when all survivors agree. When budget is exhausted with both targets present, it returns:

UNKNOWN_WITH_RESIDUAL_AMBIGUITY

The output includes the remaining changed and unchanged scenario counts, unresolved dependency alternatives, cost spent, checks skipped, and the cheapest next checks that could still separate a target-crossing pair.

Stage 6: Audit model violations and misses

An empty survivor set is a model violation, not permission to improvise. Shadow evaluation also uses hidden adjudication to identify wrong verdicts caused by actual worlds outside the model set. Those cases become candidates for adding a dependency variant, not examples to delete from the scorecard.

6. E27-A exact finite diagnostic

6.1 Scenario construction

E27-A preserves the Round 26 roots:

Root	Intended role
issuer	Origin of the claimed change record
transport	Delivery path for records
observer	Component that watches and records state
verifier	Component that checks observer outputs
external	Separately stored commitment
probe	Behavioral test path
witness	Nominally separate witness

Every scenario contains zero, one, or two faulty roots. There are $1 + 7 + 21 = 29$ possible fault sets and two target values, yielding 58 scenarios per dependency variant.

The planning model contains two variants:

1. `independent_witness`: the witness is blind only when the witness root is faulty.
2. `correlated_witness`: the witness is also blind when the issuer root is faulty.

Their union contains 116 scenarios. Evaluation adds a third variant that is not given to the planner:

3. `hidden_common_origin`: both the witness and the nominally external commitment secretly share the issuer failure root.

This third variant is a constructed counterexample, not an estimate of production frequency.

6.2 Checks and deterministic semantics

Check	Possible outputs	Simplified diagnostic semantics
receipt	change, none	Change appears unless issuer, transport, observer, or verifier is faulty
checkpoint	change, clean, missing	Transport or observer faults make it missing; issuer or verifier faults can make it falsely clean
heartbeat	ok, missing	Tests observer availability, not the target directly
transport acknowledgement	ok, missing	Tests transport availability, not the target directly
verifier canary	ok, failed	Tests verifier state, not the target directly
external commitment	change, clean	Detects change unless external root is faulty; in the hidden variant it also shares issuer failure
behavior probe	change, clean	Detects change unless probe root is faulty
independent witness	change, clean	Detects change unless witness root is faulty; correlated variants also share issuer failure

The outcome functions are deliberately simple enough to enumerate. They should be read as a test fixture for reasoning about dependencies, not a faithful production model.

6.3 Cost regimes

Check	Round 26 cost	Cheap declared witness stress
receipt	1	1
checkpoint	1	2
heartbeat	1	1
transport acknowledgement	1	1
verifier canary	1	1
external commitment	3	4
behavior probe	3	5
independent witness	3	2

These are authored cost units, not measured tokens, seconds, or dollars. The second regime creates a temptation to query a cheap witness whose independence may be overstated.

6.4 Policies

Policy	Planning model	Acquisition rule	Adjudication model
sacp_fixed	Two variants	Status-heavy fixed order	116-scenario gate
target_first_fixed	Two variants	Target-facing fixed order	116-scenario gate
nominal_world_entropy	Independent witness only	Greedy world-entropy reduction per cost	58-scenario gate
nominal_target_entropy	Independent witness only	Greedy target-entropy reduction per cost	58-scenario gate

Policy	Planning model	Acquisition rule	Adjudication model
nominal_edge_greedy	Independent witness only	Greedy expected cross-target edge reduction per cost	58-scenario gate
safe_minimax_greedy	Two variants	Greedy worst-branch cross-target pair reduction per cost	116-scenario gate
safe_uniform_dp	Two variants	Exact finite-horizon expected categorical resolution	116-scenario gate

The fixed orders are important baselines. A complex scheduler should not receive credit for beating only a deliberately weak status-first sequence.

6.5 Budgets, evaluation weights, and total runs

Budgets are 3, 5, 7, and 9 units. Every policy receives the same cap within a cell. Policies may stop early after a categorical verdict or when no remaining check fits.

The same exact policy-world traces are scored under three weighting distributions:

- **uniform:** all 58 worlds in an actual variant receive equal weight;
- **benign:** quiet worlds and smaller fault sets receive more weight;
- **issuer shift:** changed worlds receive 0.7 total mass, quiet worlds 0.3, larger fault sets are more common, and fault sets containing the issuer receive an 8x multiplier before normalization.

The experiment contains 2 cost regimes x 3 actual variants x 4 budgets x 7 policies x 58 worlds = 9,744 policy-world runs. Reweighting those traces under three distributions yields 504 aggregate cells. Reweighting does not create new observations.

6.6 Metrics

The primary harm metric is the harmful false-clean rate conditional on actual change. Other metrics are:

- correct categorical resolution rate over all worlds;
- abstention rate;
- false-change rate conditional on quiet worlds;
- separate changed-world and quiet-world resolution;
- mean cost used;
- residual cross-target pairs at stop;
- model-violation rate when no modeled scenario survives.

Resolution without correctness is never reported as success. A policy that decides every case by guessing NO_CHANGE has high coverage and unacceptable harm.

7. Results

7.1 In-model safety result

Across the two planning variants, both cost regimes, four budgets, 58 worlds per variant, and four policies that use the 116-scenario gate, there were 3,712 policy-world runs. None produced an incorrect categorical verdict. This is an exhaustive check of the finite diagnostic and an instance of Proposition 1, not an empirical estimate of real reliability.

The narrower nominal policies were also correct in the independent-witness variant, as expected. Errors appeared when the actual dependency graph was absent from their 58-scenario model.

7.2 Correct resolution under the uniform in-model distribution

Percent of the 116 modeled scenarios receiving a correct categorical verdict:

Cost regime	Budget	Status-first fixed	Target-first fixed	Safe myopic greedy	Safe uniform DP
Round 26	3	12.07%	12.07%	12.07%	37.93%
Round 26	5	12.07%	43.10%	34.48%	46.55%
Round 26	7	12.07%	100.00%	48.28%	100.00%
Round 26	9	48.28%	100.00%	100.00%	100.00%
Cheap witness	3	12.07%	12.07%	9.48%	37.93%
Cheap witness	5	12.07%	37.93%	40.52%	43.10%
Cheap witness	7	12.07%	43.10%	54.31%	59.48%
Cheap witness	9	43.10%	71.55%	76.72%	80.17%

The dynamic program dominates the three safe baselines under the uniform planning distribution because it optimizes that exact objective over the same action set. This is an internal optimality result, not evidence about a different real distribution.

The table also contains the negative result. At Round 26 cost and budget 5, the safe myopic greedy policy resolves 34.48 percent while the fixed target-first schedule resolves 43.10 percent. At cheap-witness cost and budget 3, myopic greedy resolves 9.48 percent while both fixed schedules resolve 12.07 percent. Conditioning actions on observations is not enough. A short-sighted objective can spend budget on checks that leave no affordable completing path.

7.3 Dependency misspecification creates false confidence

The clearest in-model-versus-nominal comparison uses the cheap-witness regime, budget 7, the correlated-witness actual variant, and the issuer-shifted weighting.

Policy	Correct categorical resolution	Abstention	False clean among changed worlds	Mean cost
Status-first fixed	4.75%	95.25%	0.00%	5.76
Target-first fixed	59.19%	40.81%	0.00%	6.22
Nominal world entropy	16.57%	83.43%	0.00%	6.66
Nominal target entropy	92.14%	0.00%	11.23%	6.26
Nominal edge greedy	92.14%	0.00%	11.23%	6.26
Safe myopic greedy	30.29%	69.71%	0.00%	5.82
Safe uniform DP	66.88%	33.12%	0.00%	6.49

The nominal target-focused policies appear attractive if one reports only their 100 percent categorical coverage. Once correctness and conditional harm are separated, their tradeoff is clear. They decide

every case, but 11.23 percent of actual changed cases are falsely labeled NO_CHANGE. The model-set gate trades coverage for safety inside its declared variants.

7.4 Worked example: a cheap witness and a shared issuer

Consider the actual world:

- dependency variant: `correlated_witness`;
- target: `CHANGE`;
- faulty roots: `issuer` and `external`;
- cost regime: cheap witness;
- budget: 7.

The nominal target-entropy policy assumes the witness is independent. It performs:

1. `independent_witness` at cost 2, observes `clean` because the witness secretly shares the faulty issuer;
2. `receipt` at cost 1, observes `none` because the issuer is faulty;
3. `external_commitment` at cost 4, observes `clean` because the external root is faulty.

Its 58-scenario nominal model has 29 surviving scenarios, all quiet, so it emits NO_CHANGE. The verdict is wrong.

The exact MSAV policy performs:

1. `heartbeat` at cost 1, observes `ok`;
2. `independent_witness` at cost 2, observes `clean`;
3. `external_commitment` at cost 4, observes `clean`.

Its 116-scenario model retains 47 scenarios and 132 changed-quiet cross-target pairs. It emits UNKNOWN_WITH_RESIDUAL_AMBIGUITY. The check sequence did not discover the change, but the gate prevented the absence of detection from becoming an assertion of stability.

The target-first fixed schedule also abstains in this world. It observes a clean checkpoint, a clean external commitment, and no receipt, leaving 36 scenarios and 128 cross-target pairs. This example shows why the paper separates schedule quality from verdict safety. The exact policy is not always the only safe response, and abstention is the correct scored outcome here.

7.5 Out-of-model common origin breaks the guarantee

The `hidden_common_origin` variant makes both the witness and external commitment share the issuer root, but the planning model includes only witness-issuer correlation. Under the issuer-shifted distribution:

Cost, budget	Policy	Correct categorical resolution	Abstention	False clean among changed worlds
Round 26, 7	Safe myopic greedy	16.57%	83.43%	0.00%
Round 26, 7	Safe uniform DP	92.14%	0.00%	11.23%
Round 26, 9	Safe myopic greedy	92.14%	0.00%	11.23%
Round 26, 9	Safe uniform DP	84.28%	0.00%	22.46%
Cheap witness, 7	Safe myopic greedy	22.43%	69.71%	11.23%
Cheap witness, 7	Safe uniform DP	22.99%	69.15%	11.23%

Cost, budget	Policy	Correct categorical resolution	Abstention	False clean among changed worlds
Cheap witness, 9	Safe myopic greedy	34.15%	42.27%	33.68%
Cheap witness, 9	Safe uniform DP	34.15%	42.27%	33.68%

More budget does not monotonically improve safety under misspecification. In the cheap-witness stress, the exact policy’s conditional false-clean rate rises from 11.23 percent at budget 7 to 33.68 percent at budget 9. Additional mutually dependent observations can eliminate modeled changed worlds and make the wrong target unanimous.

This result rules out a tempting but invalid inference: “If each extra check has positive expected information under the model, running more of them cannot hurt.” It can hurt when the model is structurally wrong.

7.6 What E27-A establishes and what it does not

E27-A establishes five facts about its own finite construction:

1. The model-set unanimity gate has zero categorical error when the actual scenario is inside the declared model and semantics.
2. A narrow nominal model can gain coverage while producing harmful false-clean verdicts under a plausible modeled correlation shift.
3. Exact finite-horizon planning can improve correct resolution over safe fixed baselines under its planning prior.
4. Myopic ambiguity reduction can underperform a simple target-first schedule.
5. Structural model omission can make more verification increase false confidence.

It does not establish the frequency of any failure in production, the best real cost model, the right prior, the completeness of the two dependency variants, or an improvement in Enceavra answer quality.

8. Claim and evidence table

Claim	Evidence status	Supporting material	Boundary
Sequential checks should be selected by downstream decision value and cost	Established general principle	S01, S03	Does not specify Enceavra’s model or loss
Greedy adaptive selection has guarantees only under structural conditions	Established with conditions	S02	Conditions are not proved for MSAV’s objective
Correlated tests can defeat common entropy heuristics	Established in adjacent formal work	S04	Different task and noise model
Experiment rankings can be prior-sensitive	Established	S06	Prior ambiguity does not cover omitted graphs
More test-time compute or review is not uniformly beneficial	Mixed, context-specific empirical evidence	S08, S09	Different models and benchmarks
A 116-scenario unanimity gate makes no categorical errors inside E27-A’s model	Exact finite result	E27-A code and tests	Only declared deterministic worlds

Claim	Evidence status	Supporting material	Boundary
Nominal target entropy produces 11.23% false clean among changed correlated worlds in one stress cell	Exact finite result	E27-A policy metrics	Authored issuer-shift weighting and costs
Safe uniform DP reaches 66.88% correct resolution with zero modeled error in that cell	Exact finite result	E27-A policy metrics	Not a production estimate
More budget can increase false confidence under omitted common origin	Exact constructed counterexample	E27-A hidden variant	Demonstrates possibility, not prevalence
MSAV will improve real Enceavra reliability at fixed compute	Open Enceavra hypothesis	Proposed E27-B	Requires hidden-adjudication shadow trial
Brain-like control allocation is the mechanism implemented by MSAV	Not claimed	S12-S14 used only as analogy	No literal brain-copying inference

9. Brain-inspired hypothesis, kept testable

The useful analogy is control allocation, not anatomy. The Expected Value of Control theory proposes that human control allocation depends on anticipated payoff, efficacy, and effort cost [S12]. Human behavioral and EEG experiments report that expected reward and efficacy separately influence invested control [S13]. Ho and colleagues model people as planning how to allocate limited information processing and test quantitative predictions in two experiments [S14].

MSAV borrows one design question from this literature: should a verifier spend effort where control is both valuable and likely to work? It does not claim that a dynamic program is a dorsal anterior cingulate cortex, that a check is a neuron, or that human cognitive control is optimal.

The analogy becomes scientific only through an ablation. A real shadow trial can compare:

- a reward-only scheduler that values the importance of a target but ignores estimated check efficacy;
- an efficacy-only scheduler that favors checks expected to return clean signals but ignores downstream harm;
- a cost-blind scheduler;
- the full cost-benefit scheduler.

If efficacy conditioning does not improve risk-coverage or cost-to-correct-resolution under distribution shift, the brain-inspired design hypothesis is weakened. The analogy contributes no evidence by itself.

10. Alternative explanations and rival mechanisms

10.1 One check is simply dominant

The dynamic program may look intelligent because the external commitment or behavior probe is strong in the authored model. The fixed target-first baseline tests this explanation. At Round 26 cost and budget 7 it ties the exact policy at 100 percent modeled resolution, showing that adaptivity adds no benefit in that cell.

10.2 The prior creates the result

The exact policy is optimized for a uniform planning distribution. Different weights can favor a different tree. E27-A therefore reports uniform, benign, and issuer-shifted evaluations. A real trial must fit planning weights on one partition and evaluate on a locked, shifted partition.

10.3 The cost table is handpicked

Both cost regimes are authored. The cheap-witness regime intentionally creates a difficult tradeoff. Real costs must include latency, token use, external requests, human review, and opportunity cost, then be frozen before evaluation.

10.4 Deterministic observations make unanimity too easy

Yes. Real checks can be noisy, delayed, missing, or strategically manipulated. A probabilistic extension would need a calibrated selective decision rule, not literal set elimination. E27-A is a boundary-finding diagnostic.

10.5 The model set is just a larger prior

No. A mixture over known graphs changes weights inside the set. It still cannot represent an omitted graph. The hidden-common-origin result distinguishes distributional robustness from structural coverage.

10.6 Abstention hides poor performance

Abstention is reported, not discarded. A policy that avoids error by always abstaining scores zero correct resolution and 100 percent abstention. Risk-coverage curves, changed-world coverage, and budget-normalized utility are required alongside harm rates.

10.7 A learned verifier could replace explicit lineage

A learned verifier may improve prediction, but if its training data, retrieval path, and upstream issuer share failures, a high score can reproduce common-origin confidence. Learned verification should be evaluated as another evidence channel with a lineage claim, not as an exception to lineage accounting.

11. Falsifiable hypotheses

H27.1: Conditional safety

For every scenario inside a correctly specified finite model set, a target-unanimity gate will never emit the wrong categorical target.

Falsifier: an executable in-model trace that retains the actual scenario but emits the opposite target. In deterministic E27-A this would indicate a code or specification error.

H27.2: Dependency-aware gating reduces harmful false clean

At matched compute in strata with correlated evidence roots, a model-set gate will have a lower harmful false-clean rate than a nominal independence gate, while retaining nonzero correct resolution.

Falsifier: no reduction in false clean, or a reduction achieved only by near-total abstention beyond the preregistered utility boundary.

H27.3: Non-myopic allocation improves modeled resolution

Given the same declared model, costs, and planning distribution, finite-horizon optimization will weakly dominate any included fixed schedule on expected categorical resolution.

Falsifier: a fixed schedule with higher expected resolution under exactly the same state, action, cost, and prior definitions. E27-A tests this exhaustively for the included baselines.

H27.4: Myopic adaptive allocation is not reliably better than fixed allocation

A greedy next-check policy can underperform a target-first fixed schedule because an attractive immediate reduction can block an affordable completing path.

Falsifier: across a preregistered benchmark family satisfying the claimed range of costs and dependencies, the greedy policy never loses and the difference cannot be attributed to low power.

H27.5: Structural omission can make extra verification harmful

When apparently independent checks share an omitted root, increasing budget can increase the conditional false-clean rate.

Falsifier: under hidden-adjudication common-origin strata, false confidence never increases with budget after controlling for policy and coverage, across a sufficiently powered locked evaluation.

H27.6: Efficacy-aware allocation adds value beyond reward and cost alone

A scheduler that models both target value and check efficacy will improve cost-to-correct-resolution or selective risk under shifts in check reliability.

Falsifier: efficacy conditioning provides no improvement over reward-only and cost-only ablations, or degrades worst-stratum harm.

H27.7: Residual ambiguity predicts future corrections

Among abstained cases, the recorded number and structure of cross-target scenario pairs will predict which additional checks later change the answer.

Falsifier: residual ambiguity features do not outperform simple uncertainty or random rankings on locked future-correction prediction.

12. Proposed real experiment E27-B

12.1 Purpose and safety posture

E27-B should be a shadow experiment only. It must not suppress updates, change user-facing answers, publish claims, or authorize deployment. Its purpose is to test whether the allocation and gating ideas survive natural noise and real lineage.

12.2 Corpus and adjudication

Construct a locked corpus with natural and sandbox strata reported separately:

- natural multi-authority change records collected during an observation window;
- quiet controls where hidden adjudication verifies no relevant change;
- sandbox changes injected at known times into isolated test authorities;
- common-origin cases where two visible channels intentionally share an upstream source;
- transport, observer, and verifier faults introduced independently where operationally safe.

An adjudication team or independently rooted system should determine the target without seeing the policy verdict. Cases without adequate independent adjudication remain unscored and are reported as such.

12.3 Arms

At minimum, compare:

1. status-first fixed schedule;
2. target-first fixed schedule;
3. nominal target-entropy policy;
4. dependency-aware myopic policy;
5. model-set adaptive policy optimized on a training partition;
6. random affordable check selection;
7. oracle-cost retrospective upper bound, clearly labeled non-deployable.

The oracle may use hidden labels only after collection to estimate headroom. It must never choose live actions.

12.4 Controlled total compute

All online arms receive the same per-case maximum budget, maximum external requests, token envelope, and wall-clock deadline. Early stopping is allowed and its unused cost is reported, not silently spent on extra candidate generation. Generator outputs, starting evidence, and task order are held constant across paired arms wherever feasible.

Primary comparison is at equal budget cap. Secondary comparisons report cost to first correct categorical verdict and area under the selective risk-coverage curve. A policy cannot claim efficiency by receiving a larger hidden context or a more capable verifier.

12.5 Baselines and ablations

Required ablations are:

- remove the alternate dependency graph;
- replace target-relative value with world entropy;
- remove conditioning on observed outcomes and use the induced fixed order;
- remove cost normalization;
- replace learned efficacy with constant efficacy;
- remove the unanimity gate and force a verdict;
- replace the model set with a single best-fit graph;
- hide lineage labels while keeping outcomes unchanged.

These ablations distinguish value from complexity, adaptivity, lineage representation, and abstention.

12.6 Metrics and regression detection

Report all metrics overall and by target, source authority, dependency lineage, fault type, task difficulty, and natural versus sandbox stratum:

- harmful false-clean rate conditional on change;
- false-change rate conditional on quiet;
- correct categorical resolution rate;
- abstention and model-violation rates;
- selective risk-coverage curve;

- Brier score and expected calibration error for any probabilistic output;
- mean and tail cost, latency, tokens, and external calls;
- cost to first correct resolution;
- time to restored correctness after a detected model violation;
- lineage correction rate after hidden adjudication;
- worst-stratum harm and the gap between best and worst strata.

Use paired bootstrap intervals for rate and cost differences when cases are paired. Use exact or score intervals for rare harms, and publish counts with every percentage. Predefine a noninferiority boundary for false clean and a minimum useful coverage gain before unblinding. A policy fails the gate if its overall average improves while a preregistered high-harm stratum regresses beyond the boundary.

12.7 Decision rule

Do not recommend deployment from one shadow trial. Advance only if:

- harmful false clean is noninferior overall and in every protected failure stratum;
- at least one useful efficiency or correct-resolution endpoint improves at fixed budget;
- model violations are surfaced rather than forced into categorical answers;
- results reproduce on a later time split with frozen policy and costs;
- natural and sandbox results are not combined to hide a failure in either.

13. Expected failure modes

1. **Unmodeled common origin:** two checks share a source that the dependency manifest marks as separate.
2. **Prior drift:** faults and changes move away from the planning distribution, causing a poor but still model-safe schedule.
3. **Cost drift:** a nominally cheap check becomes slow or expensive and crowds out a completing path.
4. **Noisy semantics:** a check sometimes returns an outcome not represented in the deterministic model.
5. **Policy overfitting:** the scheduler learns quirks of the development corpus and loses on a future split.
6. **Verifier gaming:** candidate agents learn which evidence pattern triggers a categorical gate.
7. **Abstention collapse:** a conservative model set grows until useful coverage approaches zero.
8. **Confidence laundering:** a categorical verdict is shown without its model-set boundary and residual assumptions.
9. **Budget laundering:** an adaptive arm receives extra context, retries, or stronger tools outside the accounted budget.
10. **Delayed evidence:** the best check arrives after the decision deadline, so nominal value is not operational value.
11. **Correlated adjudication:** the hidden judge shares the same source as the policies it evaluates.
12. **More-evidence harm:** extra dependent checks make the wrong modeled target unanimous.

14. Limitations

The finite model is small, deterministic, binary, and hand-authored. Its seven roots and eight checks are abstractions. Faults do not have temporal duration, strategic behavior, graded corruption, or partial recovery. Costs are integers without measured variance. The exact dynamic program is exponential in the remaining checks and belief states and is included as a reference policy, not a production implementation.

The uniform planning prior is convenient rather than empirically justified. The benign and issuer-shift distributions are also authored. The hidden-common-origin stress demonstrates possibility but

does not estimate prevalence. Because every scenario is enumerated, confidence intervals would be inappropriate for E27-A; the reported percentages are exact weighted results for the constructed model.

The literature is adjacent rather than directly validating. Active learning, Bayesian experimental design, sequential testing, test-time scaling, abstention, and cognitive control each illuminate part of the problem, but none provides an off-the-shelf guarantee for Encevera collaboration.

The source search was broad but not a formal systematic review. The October 14, 2026 budgeted-verification paper is a recent preprint and may change. No claim in this round depends on accepting its empirical results.

Most importantly, the conditional safety theorem can invite overconfidence if its condition is dropped. The diagnostic shows this failure directly. A model-set gate is only as complete as the dependency alternatives it contains and only as faithful as its observation semantics.

15. Practical recommendations

1. Preserve UNKNOWN as a first-class result. Never convert budget exhaustion into NO_CHANGE.
2. Keep evidence lineage separate from evidence content. Two agreeing strings can have one root.
3. Declare the acquisition objective. “Information gain” is incomplete without saying information about what.
4. Benchmark against a target-first fixed schedule, not only a weak status-first schedule.
5. Separate the decision gate from the scheduling heuristic so a cheaper planner cannot weaken adjudication silently.
6. Report conditional harm and coverage together. A 100 percent resolution rate can conceal false confidence.
7. Stress omitted graphs. Robustness over weights is not robustness to a missing dependency structure.
8. Treat extra verification as an intervention that can harm under misspecification, not as a monotone good.
9. Freeze costs and model alternatives before locked evaluation.
10. Do not deploy from E27-A. Run E27-B in shadow mode with independent adjudication.

16. Next-round handoff

The preferred Round 28 task is empirical, not another mechanism paper:

Run a fixed-compute shadow comparison of model-set adaptive verification, target-first fixed verification, and nominal target-entropy verification over independently adjudicated natural and sandbox strata. Freeze lineage graphs, costs, and policies before evaluation. Publish every wrong categorical verdict, abstention, model violation, root-lineage correction, and cost overrun.

If independent adjudication or a genuinely multi-authority observation window is unavailable, do not simulate a production result. Build a trace-replay benchmark with withheld dependency graphs and delayed reveal of the adjudicated target. The key Round 28 question should be whether residual ambiguity and explicit model violations predict real corrections better than ordinary confidence.

17. Reproducibility

The bundle contains:

- this manuscript in Markdown and PDF;
- round27_budgeted_verification.py;
- test_round27_budgeted_verification.py;
- exact CSV outputs for policy metrics, first actions, and worked traces;
- summary.json;
- the source-inspection ledger;

- checksums and run instructions.

The Python program uses only the standard library. It generates every aggregate table from exhaustive policy-world traces. Six deterministic unit tests verify scenario counts, in-model categorical safety, a nominal false-clean counterexample, an out-of-model false-clean counterexample, exact-policy dominance over included safe baselines under the uniform planning objective, and deterministic artifact generation.

18. Sources

[S01] Russell, S., and Wefald, E. (1991). “Principles of metareasoning.” *Artificial Intelligence*. May 1991. [https://doi.org/10.1016/0004-3702\(91\)90015-C](https://doi.org/10.1016/0004-3702(91)90015-C). Inspected: publisher abstract and bibliographic record. Supports value-based selection of computational actions.

[S02] Golovin, D., and Krause, A. (2011). “Adaptive Submodularity: Theory and Applications in Active Learning and Stochastic Optimization.” *Journal of Artificial Intelligence Research*. <https://arxiv.org/pdf/1003.3967>. Inspected: full 60-page paper, including formal policy definitions, cost treatment, active-learning results, assumptions, and discussion of incorrect priors.

[S03] Hay, N., Russell, S., Tolpin, D., and Shimony, S. E. (2012). “Selecting Computations: Theory and Applications.” *UAI 2012, PMLR R10*, 344-353. <https://proceedings.mlr.press/r10/hay12a.html>. Inspected: proceedings page, abstract, metadata, and paper record.

[S04] Chen, Y., Hassani, S. H., and Krause, A. (2017). “Near-optimal Bayesian Active Learning with Correlated and Noisy Tests.” *AISTATS 2017, PMLR 54*, 223-231. <https://proceedings.mlr.press/v54/chen17b/chen17b.pdf>. Inspected: full nine-page paper.

[S05] Kirsch, A., Rainforth, T., and Gal, Y. (2021). “Test Distribution-Aware Active Learning: A Principled Approach Against Distribution Shift and Outliers.” <https://arxiv.org/pdf/2106.11719>. Inspected: full 14-page preliminary manuscript.

[S06] Go, J., and Isaac, T. (2022). “Robust Expected Information Gain for Optimal Bayesian Experimental Design Using Ambiguity Sets.” *UAI 2022, PMLR 180*, 728-737. <https://proceedings.mlr.press/v180/go22a/go22a.pdf>. Inspected: full ten-page paper.

[S07] Vershinin, G., Cohen, A., and Gurewitz, O. (2025). “Active Sequential Hypothesis Testing with Non-Homogeneous Costs.” Preprint. <https://arxiv.org/pdf/2509.11632>. Inspected: full five-page preprint. Status: not peer-reviewed in the inspected record.

[S08] Bilal, A., Mohsin, M. A., Umer, M., Subhan, A., Rizwan, H., Mohsin, A., and Hougen, D. F. (2026). “What If We Allocate Test-Time Compute Adaptively?” *ICML 2026, PMLR 306*, 8277-8301. <https://proceedings.mlr.press/v306/bilal26a.html>. Inspected: official proceedings page, abstract, metadata, and linked paper record.

[S09] Feng, Y., Kempe, J., Zhang, C., Jain, P., and Hartshorn, A. (2026). “What Characterizes Effective Reasoning? Revisiting Length, Review, and Structure of CoT.” *ICML 2026, PMLR 306*, 30066-30087. <https://proceedings.mlr.press/v306/feng26e.html>. Inspected: official proceedings page, abstract, metadata, and linked paper record.

[S10] Xue, B., Cheng, J., Lyu, S.-H., Wan, Y., and Qiu, S. (2026). “Test-Time Scaling via Budgeted Multi-Attribute Verification.” Version dated October 14, 2026. <https://arxiv.org/pdf/2609.34322>. Inspected: full 24-page preprint. Status: very recent preprint.

[S11] Yu, J., and Blanchard, M. (2026). “Distribution-Free Sequential Prediction with Abstentions.” *COLT 2026, PMLR 336*, 6976-7011. <https://proceedings.mlr.press/v336/yu26a.html>. Inspected: official proceedings page, abstract, metadata, and linked paper record.

[S12] Shenhav, A., Botvinick, M. M., and Cohen, J. D. (2013). “The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function.” *Neuron* 79(2), 217-240.

<https://pubmed.ncbi.nlm.nih.gov/23889930/>. Inspected: PubMed record and full PMC theory article. Used only to derive a testable cost-benefit analogy.

[S13] Fromer, R., Lin, H., Dean Wolf, C. K., Inzlicht, M., and Shenhav, A. (2021). “Expectations of Reward and Efficacy Guide Cognitive Control Allocation.” *Nature Communications* 12, 1030. <https://pmc.ncbi.nlm.nih.gov/articles/PMC7884731/>. Inspected: full primary article, including methods, results, and discussion.

[S14] Ho, M. K., Abel, D., Cohen, J. D., Littman, M. L., and Griffiths, T. L. (2020). “The Efficiency of Human Cognition Reflects Planned Information Processing.” *AAAI 2020*, 1300-1307. <https://ojs.aaai.org/index.php/AAAI/article/view/5485>. Inspected: official paper page and full linked paper, including formal objective and two human experiments.

19. Final conclusion

Round 27 narrows the next design move. Encevera should not allocate verification by raw confidence, generic information, or the number of agreeing agents. It should represent plausible dependency structures, value checks against the decision target, and require surviving model agreement before making a categorical claim.

The exact diagnostic supports that architecture inside its declared model and exposes two limits. First, a myopic adaptive scheduler can be worse than a simple fixed schedule. Second, a model-set gate can still become confidently wrong when the actual common origin is missing from the set. More budget can amplify that error.

The practical lesson is not that MSAV is proven safe. It is that allocation and adjudication must be evaluated separately, and every safety statement must name the model set on which it depends. The next credible step is a hidden-adjudication, fixed-compute shadow trial that treats abstention, false stability, and lineage correction as first-class outcomes.