Liva.Fuxi: Worlds Change, Consequences Continue

Supplementary Methods and Evidence

Huaiyang Liu DeepBean Limited
wyonliu@gmail.com

21 November 2026

This supplement presents formal conditions, experimental methods, negative results and audit corrections. Cross-references to numbered main sections refer to the main manuscript. Experimental paths are relative to the accompanying evidence archive. Main Section 5 presents the complete-activity results.

A Additional mathematical conditions

The following standard or restricted conditions are retained as supporting analysis, not claimed as new general theory.

A.1 Finite influence under declared locality

Proposition 1, finite causal influence in a local discrete model. Assume synchronous updates on a fixed dependency graph; each component at tick t+1 reads only its declared incoming neighbors and its own state at tick t; exogenous inputs outside an intervention remain equal. Then an intervention at a component at tick 0 cannot affect a component at directed graph distance greater than k by tick k.

Proof. At tick 0 only intervened components differ. If a component's parents agree at tick k, its update and exogenous input agree at tick k+1. Induction confines differences to successive outgoing neighborhoods. □

This is an elementary locality property, not a new causal-discovery theorem. A global clock, broadcast, resource ledger, or new interaction edge may make the neighborhood global. The five-phase witness has explicit within-tick dependencies, so the proposition applies to its unrolled phase graph, not indiscriminately to one entire tick as a single edge. It motivates affected-region simulation only when the actual dependency structure justifies it.

A.2 Legacy-path simulation

An extension contains new state and action definitions, a transition update, a migration mapping, preserved properties, and an activation scope. Existing capabilities may first be composed. New primitives require an executable definition and an isolation boundary. Changes to existing institutions or physics require an explicit account of affected behavior; they cannot automatically claim conservative compatibility.

Proposition 2, legacy-path preservation. Let m be an injective mapping from old states to new states. Assume correct initial mapping and, for every reachable old state s and enabled old action u, that u remains enabled at m(s) and

$$T'(m(s), u) = m(T(s, u)). \tag{S1}$$

Then every finite old action trace has a corresponding mapped trace in the new system.

Proof. The empty trace follows from the initial mapping. If a prefix maps correctly, enablement preservation permits its next action, and the transition equality maps the successor. Induction completes the argument.

This is a standard sufficient simulation condition, not a new general theorem. It does not imply that arbitrary new actions preserve every old continuation, nor does it establish liveness or semantic equivalence of physical processes. Stochastic systems require an appropriate pushforward relation between transition kernels.

The intended learning outcome goes beyond this proposition. Given a participant's information, capabilities, and budget H, let R contain classes of outcomes that can actually be reached. An extension can add reachable outcomes while preserving specified old behavior. Counting states is not enough: outcomes require a declared equivalence relation, and practical use requires discovery and competence. Human significance cannot be inferred from state cardinality.

The empirical hypothesis therefore concerns the adoption and reuse of valid new mechanisms by unfamiliar participants, under matched access and compute. Creation quality, preservation of earlier consequences, and transfer are measured separately.

A.3 Restricted result reuse, including RHS dependencies

Consider a finite deterministic DSL in which action a has an explicit guard reading G_a, one assignment s_j := f_a(s_Ra), and a positive constant cost c_a. The finite sets G_a and R_a include all guard and right-hand-side reads. For a goal predicate ϕ, initialize D with its read variables. Whenever an old or new action writes a variable in D, add both G_a and R_a; repeat to a fixed point. Construct the closure using the union of the old and new rules, including removed and added writers. State-dependent costs, resource checks or implicit reads must also enter the closure or fall outside this restricted statement. Unknown effects cannot be declared independent.

Proposition 3, corrected contract-result reuse in the restricted DSL. Suppose the goal and upper cost budget are unchanged, the old result is validated, the two initial states agree on D, and the sets of actions writing into D agree in guards, assignment functions and constant costs. Suppose outside actions cannot affect time constraints, observations, scheduling or resource availability relevant to the contract. Under these conditions, the bounded existential reachability result is unchanged.

Proof. Every relevant guard and assignment value is a function of the projected state because closure includes G_a and R_a. Paired executions of a relevant action therefore have identical enabledness, projected successors and costs. Outside actions leave this projection unchanged and have positive cost; removing them cannot invalidate a witnessing path with an upper cost bound under the stated absence of implicit dependencies. Every old witnessing path thus projects to a new path and conversely. All goal variables are included, giving equality of the Boolean results. □

An earlier formulation included guards but omitted right-hand-side reads and was false for general unary assignments. For example, y := x has no guard. With goal y=1, budget one, and initial y=0, changing x from one to zero changes reachability even though the guard-only projection D={y} is unchanged. The corrected closure adds x. The older executable extension suite used constant assignments, whose right-hand sides read no variables; its recorded results survive this correction. Supplement C.6 adds an explicit read-dependent diagnostic rather than treating the old suite as evidence for the broader claim.

This is a domain-restricted consequence of established dependency analysis, projection, and incremental verification, not a new general theorem. It does not cover mandatory external events, fairness, observational knowledge, adversarial scheduling, or arbitrary plugin code. In those settings, assumptions must be represented and invalidated, or a stronger verifier must be used. Reuse of a stochastic certificate additionally depends on the transition kernel and probability bound, not just reachable states.

A production cache key would bind contract semantics, world and branch, projected initial state, rule and observation versions, environment assumptions, remaining budget, and verifier version. Changes in time or participant permissions can invalidate a certificate without any visible geometry change. Dependency metadata should be derived from restricted executable syntax or checked instrumentation. A plugin's self-declared footprint is not sufficient evidence.

B Further analytical conditions

Proposition 4, finite-horizon approximation. Assume identical initial distributions, the same fixed policy, a common Markov state sufficient for both transition kernels, the same reward function bounded in [0,1], and uniform total-variation error at most ϵ between transition kernels. A partial observation that is not a sufficient state cannot be substituted without further assumptions. For H rewards at times 0 through H−1,

$$|V_H - \widehat{V}_H| \leq \sum_{t=0}^{H-1} \min(1, t\epsilon) \leq \epsilon H(H-1)/2. \tag{S2}$$

Argument. Couple transitions maximally while trajectories agree. The probability of a mismatch by time t is at most tϵ. A reward difference is at most one after mismatch. Summing gives the bound. This familiar coupling argument motivates explicit error budgets; prediction scores on a finite dataset do not establish its uniform assumption.

Proposition 5, finite observational isolation. Suppose two initial states are indistinguishable to participant i, and each permitted step has equal conditional distributions over i-visible outputs with related successor states. Their finite observation-trace distributions are equal by induction. The premise must include timing, return codes, caches, and observable world changes. Filtering an initial prompt alone is insufficient. Authorized disclosure requires a separately specified channel.

These propositions are local obligations. They do not prove arbitrary generated programs correct, establish universal fairness, or certify all information channels in a deployment.

B.1 Conditional recovery properties

Proposition 6, retained finite consequence closure. Suppose each local handler atomically commits its effect record, state transition and all derived outbox effects; stable effect IDs are unique within a declared semantic namespace; the receiver commits an effect at most once; durable handoff never discards the last recoverable copy; and all necessary owners, storage and channels eventually recover and fairly serve pending work. For an admitted root whose derived effect closure is finite, every effect in that closure eventually commits exactly once in its authority scope.

Proof. A committed parent leaves each derived child in recoverable storage. Eventual service and recovery deliver that child to its owner. The owner either commits it atomically or finds its prior commit. Induction over derivation depth establishes eventual commitment; the finite closure excludes infinite unfinished depth. Receiver deduplication establishes at-most-once application. This proves neither one transport delivery nor an atomic global transition. If an effect is rejected by an execution precondition, the committed outcome is that rejection, not the originally requested change. Cyclic systems require a finite unfolded effect closure or a weaker prefix property. External non-idempotent actions and loss of every durable replica violate the premises. □

This is an application of established transactional messaging reasoning [37], not a new exactly-once algorithm. The local experiment assumes coordinator survival, and therefore cannot establish the full production premise about durable handoff.

Proposition 7, reserved-resource conservation. For a resource owner with initial stock B, let initially nonnegative available quantity A, live reservations R and consumed quantity C satisfy A+R+C=B. For q≥0, serial atomic transitions reserve q by moving q from A to R when A≥q, consume q by moving q from R to C under a unique token, and release q by moving q from R to A. If tokens cannot be consumed or released twice, all three quantities remain nonnegative and the equality is invariant.

Proof. Each transition preserves the sum and tests sufficient quantity in its source term. Token state makes replay a no-op. Induction over the owner's committed order gives the claim. Independent owners cannot use this proof to guarantee a multi-resource transaction without an additional protocol; expiring and consuming a token must share the same authoritative order. □

The experiment implements the simpler immediate allocation A+C=B, not the entire reservation/expiry protocol. Without deduplication, two accepted attempts for one logical task could consume twice while retaining nonnegative stock. This is a possible counterexample, not the observed failure in Supplement C.5: the recorded schedules instead produced inconsistent task rejection after inventory exhaustion.

Proposition 8, overload boundary. Assume at most μ units of admitted effect work can finish per unit time and lossless execution. If $A(t)-\mu t$ is unbounded above, no implementation can keep outstanding admitted work uniformly bounded. This follows directly from the accounting inequality in main Section 3.3. The proposition says nothing about whether a chosen approximation is acceptable; changing the admissible world semantics must be reported and evaluated separately.

B.2 External service obligations and end-to-end evidence

Resource safety alone permits rejecting every request. For the experiments below, the external workload specifies R distinct authorized, unit-cost requests, initial capacity B, no cancellation or expiration, eventual recovery and service, and a complete drain. Its minimum required number of fulfilled requests is K=min(R,B). This requirement is fixed before observing grants; the legal winning subset may vary with scheduling. The scorer checks the submitted set, K resource uses, the complete derived world effects, terminal states, remaining stock, and retained execution evidence. This is a bounded service agreement, not a general liveness theorem for arbitrary agents or replenishing resources.

The end-to-end obligation also requires the correct subject to observe a task whose identity, context, outcome and receipt match its confirmed intention and actual world execution. An internally completed world with clients told that all tasks were rejected is a failure. The HTTP scorer therefore independently binds each observation to its confirmation and response trace, and each success artifact to world, branch, root, semantic version, goal, committed effect and timestamp. World correctness and delivered fulfillment are reported separately. Traces are trusted experimental evidence protected by recorded hashes; these checks are not a defense against coordinated fabrication of every data source.

C Preserved experiments, costs and corrections

The accompanying standard-library program explores a three-holder, one-object atomic transition abstraction through depth five, including transfer, duplicate submission, invalid ownership, and revocation. It counts 1,365 states including logs and checks its stated local properties. This does not test a database or distributed network.

Two conflicting transfer intentions give two owners under arrival-order execution. Sorting the same admitted set by a fixed policy gives one outcome. This demonstrates a mechanism, not a universal fairness guarantee.

An authored extension graph grows from four reachable states to nine while preserving mapped legacy paths and adding a festival outcome. A destructive variant removes the old crossing operation and fails the preservation condition. No learned model generated this graph, and these counts provide no evidence of creativity or user value.

A privacy witness gives identical initial views but different result codes, demonstrating a receipt channel. An external-effect witness gives identical missing acknowledgments for executed and non-executed actions. Event-time calculations reproduce a completion probability of approximately 0.6335, refusal probability 0.3167, and deadline mass 0.0498 for the specified constant hazards and ten-second deadline. The code and raw outputs are local companions to this draft.

C.1 Qualitative causal-world witness

The second source package, `experiments/causal_world/`, implements a generic site interaction interpreter. It does not branch on "fireplace," "grass," or organism names. Every tick applies explicit

phases: rainfall changes local wetness; dry fueled sites ignite from a spark or previously burning neighbors; burning consumes fuel; a deterministic organism policy chooses a lower-risk available location from prior heat observations; exposure accumulates and may produce a cooked-object state. Two ticks of exposure triggering this state is an authored game convention, not biology or calibrated heat transfer. Wetness does not evaporate in this short model.

A factorial diagnostic varies rainfall, initial wetness, middle-site fuel, an escape route, and a thermal link, producing 32 configurations of six ticks. Eight ignite the source; two ignite the middle material; two produce escape; one produces a cooked-object state. These frequencies describe the authored grid, not natural probabilities. Exact stage traces agree with a separately expressed synchronous classical reference in every configuration. The primary simulator also agrees across 192 permutations of site iteration order; adding a disconnected fuel object leaves all 32 reported outcomes unchanged. Iteration-order invariance is weaker than invariance to renaming: movement tie-breaking still uses stable IDs.

Table S1: Qualitative causal-world conditions.

Controlled condition	**Source ignition**	**Material spread**	**Escape**	**Cooked-object state**
Dry, trapped	Yes	Yes	No	Yes
Dry, open route	Yes	Yes	Yes	No
Rain, trapped	No	No	No	No
Rain, all sites sheltered	Yes	Yes	No	Yes
No spark	No	No	No	No
Middle material lacks fuel	Yes	No	No	No
Source thermal link removed	Yes	No	No	No

All use the same local rules; no outcome-specific story event is scheduled. The model records the local inputs behind ignition, movement, and conversion events, and paired interventions expose dependence on those inputs. Removing all intersite heat coupling changes three of the 32 outcome vectors. This ablation establishes that interactions matter in these cases, not that the chosen decomposition outperforms a general engine.

The direct classical reference ties exactly, which is the expected result. No model learned the rules, inferred hidden causal structure, or autonomously generated an extension. The witness does not reproduce every narrative detail of the motivating example: the organism escapes when a lower-risk route is available and transforms only when trapped with sufficient exposure; it is not scripted to jump into fire. It makes the intended research object concrete while leaving intelligent prediction, physical fidelity, and situated discovery unproven.

C.2 Extension-contract diagnostic suite

We additionally implemented a standard-library Python interpreter and shortest-path checker for the restricted DSL. Stages 0–3 represent preparation, materials, assembly, and completion. Twelve authored configurations combine 16 or 64 components, independent/chain/fan-in dependency structure in groups of four, and two deterministic initial configurations. Each configuration receives eighteen independent proposals: six change classes at three targets. Every proposal starts from the same validated baseline. A contract's horizon is its baseline shortest completion cost; this deliberately tight budget makes time regressions observable.

The changes add a shortcut, remove assembly, replace an action identifier with an equivalent operation, reset one initial state, add a prerequisite, or double all action costs. They exercise changes in executable behavior, but do not represent unrestricted generated programs. Eighteen reset cases are no-ops because the initial value is already zero. Thirty-six global-cost proposals include duplicate

targets: global cost does not depend on target. These are reported as structured cases, not independent statistical samples.

FULL revalidates every active contract using exact projected search. REUSE uses the same checker only for potentially affected contracts. STATE-ONLY checks migrated state ranges without examining future feasibility. It is an ablation, not the strongest competitor. STRONG-CACHED is a strong extensible runtime with the same standard dependency cache; it is algorithmically identical to REUSE. Its tie is declared by construction, not claimed as an independent implementation or reproduction of INVEST.

Table S2: Finite extension-contract diagnostic.

Change scope	**Proposals**	**Accepted by FULL and REUSE**	**FULL solver calls**	**REUSE solver calls**	**STATE-ONLY false admissions**
Local changes	180	106	7,200	270	74
Global cost changes	36	0	1,440	1,440	36
Total	216	106	8,640	1,710	110

All 8,640 individual contract decisions agree between FULL and REUSE. The corresponding expanded-state counts are 83,780 and 16,986. These are deterministic operation counts, not end-to-end speedups: dependency extraction, cache construction, compiler analysis, memory, concurrent execution, and production workloads are not included. Initial certification is required by both methods and excluded from this revalidation-only table. Savings depend on modularity; global coupling removes them. STRONG-CACHED ties REUSE exactly, so this experiment provides no Fuxi-specific algorithmic advantage.

An independent search path disables projection and explores all four state variables. Across all sixteen initial 0/1 assignments, three dependency structures, four targets, and six changes, 4,608 goal comparisons match projected search and cached decisions. This checks the projection optimization against full-state search while sharing the transition interpreter; it is not an independent formal proof of the interpreter.

C.3 Counterexamples and negative findings

A dependency witness has only the second component's contract active. It requires an assembly step in the first component. Removing that step makes the active contract impossible. Deliberately omitting the cross-component read from metadata causes REUSE to keep a true certificate and admit the change. FULL rejects it. This validates the need for complete dependencies rather than reliance on author declarations.

An equivalent operation replacement removes an old action identifier but retains all specified bounded goals. Thus literal legacy action-trace preservation would reject some permissible redesigns. This witness is an action-renaming example; broader semantic replacements still need verification and it does not show changed human meaning remains acceptable.

Two further witnesses were added during the post-run scope audit and are not part of the original proposal comparison. With one unit of stock and two consumers, either promise is individually achievable but both cannot be fulfilled. With two equiprobable hidden keys and identical observations, a state-aware planner succeeds with probability one, while any observation-measurable single action succeeds with probability at most one half; randomization cannot improve that value. The finite suite's existential contracts therefore establish neither jointly satisfiable social commitments nor situated competence.

All source, protocol, per-proposal records, and code hashes are preserved under `experiments/extension_contracts/`. The previous abstraction checks remain under `formal/`. These local records contain no new training result, external engine benchmark, or human study.

C.4 Exact identification counterexamples

A new standard-library diagnostic enumerates binary structural models using rational arithmetic. It illustrates established identification limits [26, 27]; no parameters are fitted.

First, let U be a fair bit and let both models set X=U. Model A sets Y=U; model B sets Y=X. Both observational joint distributions put probability one half on each of (0,0) and (1,1), so observational prediction cannot distinguish them. Under do(X=1), however, the expected Y is one half in A and one in B. After observing X=0,Y=0, the counterfactual Y under X=1 is zero in A and one in B.

Second, let U and R be independent fair bits and X=R. Model C sets Y=U; model D sets Y=U XOR X. Both have the same uniform observational joint distribution and the same Y marginal under either intervention do(X=0) or do(X=1). Yet after observing X=0,Y=0, the counterfactual Y under X=1 is again zero in C and one in D. Thus these observational and interventional distributions do not determine the episode-specific counterfactual.

The exact outputs are preserved in `experiments/causal_identification/results.json`. They refute unrestricted transfers from prediction to intervention and from intervention marginals to unit-level counterfactuals. They do not preclude identification under additional justified assumptions, nor constitute a new causal-inference method. The diagnostic changed the manuscript's interpretation of its causal-world simulation: programmed mechanisms supply synthetic assumptions, rather than evidence that a learner discovered them.

C.5 Cross-partition execution and worker-crash recovery

The preregistered local protocol `FX-RUNTIME-04` was sealed before executing its 18 formal configurations. Development used an eight-plot chain only. Formal cases cross three topologies (24-plot chain, 31-plot binary tree, 24-plot cycle with additional edges), three request loads (4, 16 or 64 tasks, interarrival spacing 8, 2 or 0 logical units), and clean versus mixed faults. Partition counts are 3, 3 and 6 by load. A single seed fixes fault selection. These are structured stress cases, not independent samples of a deployment population or unseen learned-task generalization.

Each worker is a separate Python process with a separate SQLite database, WAL mode and FULL synchronous commits. A trusted coordinator serially drives event service over local pipes. Transport delay, duplication, reordering and a temporarily blocked link are controlled in logical time. A real worker process exits immediately after committing its first selected growth event and before replying, then restarts from its database. This exercises the lost-acknowledgment window and worker persistence. It does not inject coordinator loss, storage power failure, a real WAN partition, concurrent ownership changes, or distributed consensus faults.

All methods run identical authored rules: a task requests one scarce water unit; granted work propagates a wetting wave; a plot grows once per wave; reaching the goal plot emits a fulfillment receipt. Only half the requests can initially be granted. Each allocation is an immediate consume, not a full reservation/expiry implementation. The original final-state oracle used graph reachability and actual grants, which failed to impose a minimum service obligation. Supplement C.6 corrects this defect using external requests and initial capacity before evaluating the legitimate winning subset. Growth is a qualitative counter rather than a continuous ecological model. Cycles terminate because once-per-wave activation is part of the world rule, not because the scheduler silently drops deep events.

The candidate uses root-based scheduling, a transactional outbox, inbox deduplication and goal receipts. The strong `durable_actor` baseline has identical world semantics and reliability mechanisms with FIFO scheduling. It is a local comparator, not an Orleans or SpatialOS reproduction. Each ablation changes one candidate feature: volatile derived messages, disabled inbox deduplication, or

acceptance-based completion. No method receives a stronger model, more private information, or an extra execution tool.

Table S3: Fixed-rule partition execution and worker recovery.

Method	Correct cases	Missing plot consequences	Wrong task terminal states	Premature completions
fuxi	18/18	0	0	0
durable_actor	18/18	0	0	0
no_outbox	12/18	108	3	0
no_dedup	12/18	0	30	0
ack_only	0/18	0	0	252

All 90 runs terminated without platform error or timeout. The candidate and strong baseline both retain all passes under the corrected external-obligation audit in Supplement C.6. The no-outbox failures occur on chain/tree mixed-fault cases; redundant paths mask the lost derived message in the cyclic cases. No-dedup produces late resource rejections that overwrite working tasks in 30 instances; no duplicate consumption was observed in these particular schedules. The acceptance-based ablation completes 252 tasks before the goal is committed. These failures validate the scoped distinctions, not an original algorithmic advantage.

Round-robin scheduling does not consistently improve tail latency. For the 64-request mixed-fault chain, candidate/FIFO p99 logical terminal latency is 1241/1231; for the corresponding cycle it is 2305/2315. These old quantiles pool fulfilled and resource-rejected tasks and must not be interpreted as fulfillment speed. Supplement C.6 separates the outcomes and corrects that interpretation. These logical units are not milliseconds. Actual IPC timings, CPU, database bytes, queue maxima and every attempted delivery are reported per run. Across the 18 candidate cases, summed runner wall time is 28.82 seconds, versus 26.35 for FIFO; one unreplicated, sequentially ordered local measurement cannot establish a speed ranking. Preparation and authoring costs are not inferred from these timings.

This witness establishes more than an in-memory state check: independent worker processes actually commit persistent state and recover after exit. Its scope is still small and local. It provides no scale, long-duration, A2A wire-conformance, human-negotiation, learned-dynamics or cross-line integration result. The sealed protocol, all 90 rows, complete traces, per-partition databases, scorer, development records and environment are retained in `experiments/runtime_v04/`. There is no official benchmark score or leaderboard submission.

C.6 Review-driven corrections and retained negative results

External obligation. A reject-all mutation showed that the original scorer could pass a reject-all mutation because expected work was derived from actual grants. The corrected audit applies Supplement B.2 to all 90 retained initial runtime runs without changing their files or original scores. The case counts in Supplement C.5 remain 18/18, 18/18, 12/18, 12/18 and 0/18. Retention of those results does not vindicate the old scorer: reject-all and propagation-trace-deletion mutations passed it but fail the corrected audit. Deleting grants and skipping one plot also fail; an unmodified control passes. This is post-hoc rescoring, not 90 additional system executions.

Latency interpretation. Fulfillment and resource rejection have different meanings and must not share a headline latency. Recomputing the 36 candidate/FIFO traces gives the following successful-completion medians for four clean cases. The full audit separately retains completion and rejection p50/p95/p99 with their sample sizes for every case.

Table S4: Successful-completion latency under the corrected audit.

Case	Candidate completion p50	FIFO completion p50	Unit
tree, 4 requests	64	42	logical steps
tree, 16 requests	275	262	logical steps
chain, 16 requests	219	218	logical steps
tree, 64 requests	1136	1136	logical steps

None of these cases supports faster fulfillment. A lower pooled median can result from faster refusal while successful work remains unchanged or slower. The new study measures elapsed time and retains the outcome separation.

Dependency correction. A right-hand-side dependency counterexample exposed the missing assignment reads in Proposition 3. A diagnostic enumerates ten authored three-variable binary action sets and all 8×8 pairs of initial states per set. Among 640 comparisons, guard-only reuse incorrectly preserves 42 results; guard-and-right-hand-side closure preserves no incorrect result in this finite suite. The explicit y:=x counterexample establishes the logical defect; enumeration is a regression witness, not proof for arbitrary programs. The constant-assignment suite in Supplement C.2 and its tie with strong conventional reuse remain intact.

Client receipt correction. An internal audit found a further gap while the first HTTP study was running: replacing both client outcomes in a one-success/one-rejection fixture with rejection and removing success artifacts still passed. We retain that first 96-run series as development evidence under its original frozen scorer, without altering scores. The repaired scorer checks Supplement B.2's client binding. Fourteen deliberately corrupted client variants, including all-rejected output, wrong subject/context/task, wrong receipt root/world/branch/version/goal/time/effect, absent artifact and duplicate task, are rejected; the intact fixture passes. A separate deletion of HTTP responses fails the observation-evidence check. These controls diagnose scorer behavior rather than establish new runtime performance. The corrected formal study was frozen and rerun after this repair.

The repair source, original trace hashes, mutation outcomes and complete grouped latency records are retained in `revisions/v1.1/experiments/review_repairs/` and `revisions/v1.1/experiments/a2a_world/client-binding-controls.json`. The original runtime results and raw evidence remain unchanged.

C.7 Real HTTP cooperation and concurrent world execution

The corrected `FX-A2A-11` protocol freezes inputs, implementation and scoring before a second 96-run series. Sixteen configurations cross chain/tree topology, 16 objects, 8/32 confirmed requests, capacity at 25%/75% of requests, and clean/mixed faults. Three methods and two repeats per configuration run in alternating predeclared order. The first 96-run series remains a superseded development result after the client-scoring defect in Supplement C.6; it is not silently included in the formal denominator. Formal configurations differ from the small development chain, but no learning occurs and no unseen-task generalization is claimed.

Execution and observation. Independent clients perform real local HTTP discovery, proposal, confirmation and polling against a JSON-RPC server using A2A wire version 1.0. Up to 16 clients issue requests concurrently. Three partition processes each have a dedicated scheduling thread and SQLite database. Service-start/receive intervals show actual overlap of two or three worker RPCs in every formal run; unlike Supplement C.5, service is not globally serialized. Both root-fair and FIFO controls have identical world rules, execution capability, persistence, queues, HTTP endpoints and outcome evidence. They differ only in selecting ready work. This is a strong matched local control, not a reproduction of a mature external runtime.

The world rule consumes one water unit for a successful root and activates each reachable plot once. The root-fair policy rotates ready causal roots; FIFO selects ready arrivals. The acceptance-only ablation changes the completion predicate while keeping world execution. Requests have no learned policy or linguistic bargaining. Task completion means reaching the declared goal plot; full world correctness additionally requires all specified consequences after drain. The public observation endpoint reads partitions separately and does not provide a globally consistent snapshot.

Mixed runs inject deterministic 1–4 ms message delay and duplication, then force partition 1 to exit after a committed growth effect and before its reply. Every mixed run records one real worker exit and restart: 48 in total. A separately specified [20,70) ms blocking window after the first confirmation is schedule-dependent; 20 of 48 mixed runs actually encounter it (candidate 6, FIFO 7, ablation 7). The other 28 are not evidence of that link-block condition. The coordinator and gateway remain alive, storage is not replicated, and this is loopback fault injection rather than a WAN test.

Table S5: Concurrent HTTP execution and delivered outcomes.

Method	**Correct runs**	**Required / actual world goals**	**Correct delivered successes**	**Premature completions**
Root-fair candidate	32/32	320 / 320	320	0
Matched actor FIFO	32/32	320 / 320	320	0
Acceptance-only ablation	0/32	320 / 320	0	320

All methods ultimately execute 5,120 plot activations; only the first two deliver valid goal receipts under the full obligation. An ablation's eventual world goal cannot retroactively validate its early completion claim. All 96 formal runs terminate without platform error. The repaired binding controls independently reject missing or wrong client evidence even when the underlying world remains correct.

Outcome-separated elapsed time. Latency starts before sending confirmation and ends when polling observes the terminal state, including HTTP and polling overhead. The table pools the fixed workload mixture, with 320 observations per outcome per method. These are descriptive distributions, not paired treatment effects or independent samples from general deployments. Quantiles use floor((n−1)p); all per-case values and both repetitions are retained.

Table S6: Completion and rejection latency, reported separately.

Method and outcome	**p50 seconds**	**p95 seconds**	**p99 seconds**
Candidate completed	1.394	4.279	6.204
FIFO completed	1.788	3.338	3.520
Candidate rejected	0.115	0.648	0.912
FIFO rejected	0.304	0.842	1.004

Candidate completion p50 is lower in 15 paired runs and higher in 17. Its pooled completion tail is worse. Summed load-and-drain time is 50.259 seconds versus 47.904 for FIFO, and peak outstanding work is 143 versus 114. Per-run successful throughput ranges from 2.086 to 38.521 tasks/second for the candidate and 2.949 to 52.079 for FIFO. These ranges are not matched endpoint comparisons. No scheduling speedup or capacity advantage is established. The acceptance-only variant appears faster while violating all completion obligations, illustrating why performance must be conditioned on semantic validity.

Cost and recovery. Candidate/FIFO runner wall totals are 72.430/72.876 seconds, including startup, shutdown and audit; CPU totals are 90.857/92.960 seconds across parent and child processes. Post-drain scoring costs 1.228/2.131 seconds and is excluded from load-and-drain throughput. Each

method records 7,834 delivery attempts. Restart measurements span 0.148–0.621 seconds for the candidate and 0.117–0.487 for FIFO, including the injected 20 ms wait and process startup; they are not a production RTO guarantee. Worker RSS values are terminal samples, not aggregate peak memory. Per-run queue waits, worker RPC p50/p95/p99, pending peaks, databases and traces remain in the raw result bundle.

The 96 formal runs consume 182.627 measured runner wall seconds and 226.321 CPU seconds. The superseded 96 development runs add 189.890 wall seconds and the initial six development runs add 7.415; failed and successful integration attempts are retained separately. Runtime model tokens and external API charges are zero. Preparation-model usage, uninstrumented editing time and complete authoring cost are unknown and are not reported as zero. Every formal trace has a recorded SHA-256 hash. This is an internal protocol with no official leaderboard score.

Continuous activity and extension. The frozen integration fixture starts with eight objects and three water units. Two preassigned subjects discover, propose and confirm work, disconnect during propagation, then return to changed state and actual receipts. One worker recovers from a real post-commit exit. After drain, the fixture adds a ninth object and edge, retains existing growth, and executes a third task against the new goal: 25 required activations across the three tasks are present and all three water units are consumed. An old unconfirmed offer is rejected after the version changes; an earlier completed receipt retains its old goal and version. Duplicate confirmation, message-ID collision, context mismatch, unsupported protocol version and preconfirmation cancellation are also checked. All 18 authored integration checks pass. This is a quiescent extension, not general live migration, a generated mechanism or a participant study.

The new capability relative to our initial runtime artifact is an actual protocol-to-world-to-client path with concurrent service and a version boundary. Relative to established actor and programmable-world systems, the scientific contribution remains a proposed composition and an executable evaluation obligation, not a demonstrated algorithmic superiority. The paired negative result and the world-correct/client-wrong counterexample are substantive constraints on that proposal. The study does not test the learned consequence model, full ecology, authentic social development or humanistic value described in the broader architecture.

The complete formal HTTP dataset is at `revisions/v1.1/experiments/a2a_world/formal-v2/` in the evidence archive, beside its protocol, source hashes, comparisons and integration traces. The original runtime experiments referenced in Supplement C.1–C.5 are under `experiments/`. Artifacts in Supplement C.8–C.9 are under `next-research/2026-11-20-causal-transition/`, in the separately fixed sample-v0.2, sample-v0.3 and in-situ-v0.1 directories. Directory version labels identify preserved artifacts; they do not designate new results. The local HTTP subset is not an official registered A2A extension or full interoperability certification.

C.8 Continuation, automatic compilation and a rejected reduction hypothesis

The preserved v0.2 garden sample contains an old emitted heat pulse, a receiving material change, one shield resource, two subject-bound obligations, vegetation growth, a fleeing animal and the subsequent use of a researcher-defined instrument installed during the episode. Deadline-before, exact-deadline, late-arrival and previously-breached controls distinguish fulfillment from an endpoint that merely looks successful. A fixed shield-at-tick-3/install-at-tick-4 plan works in both wind outcomes, so this scene provides no feedback-adaptation advantage. Separate controls exercise the action-before-environment quantifier. An obligation past the horizon yields unknown; after extending the horizon, the pending harmful effect is visible.

One online fixed run snapshots at tick 2 and installs at tick 3 while the world advances. Another returns at tick 8 after ticks3–8 and the heat interaction have occurred; its plan is rejected as stale. The latter retains the old obligations but fails the expansion objective. The injected analysis delays are control conditions, not measured model inference costs. The first joint state projection remains 108 states before and after projection, with 29 visited solver states and cost2 in both same-capability methods. It has no reduction benefit.

The new v0.3 compiler uses the same scene without enlarging the load. Its declared development grid varies legal install/shield times, wind and arrival boundaries, with additional representative traces. Across 133 cases, all generated tick states agree with source interpretation and the shared fields of the independently written v0.2 runtime. Fifty cases meet the bounded goal, 82 do not and one is unknown beyond the horizon. These are exposed development cases, not held-out transfer results.

For the main model, 396 expression nodes become 260 query nodes. The generic time-aware compiler control retains exactly the same set, with the same source, constant folding and cache. Expression nodes are not the 108 game states reported above. Total fixed runner wall time is 0.930 seconds and CPU time is 0.857 seconds; six graph compilations account for about 0.120 seconds. Source interpretation, complete and reduced evaluation, reference comparison and both slice construction costs are recorded. Tiny timings are not ranked. An earlier development run unnecessarily rebuilt the generic slice per case; its source and results are preserved, and the fixed run caches both methods equally.

Binding controls reject an altered query root, missing operand edge, removed environmental input alternative, changed deadline and wrong horizon. They demonstrate the new declared trusted-construction boundary, not external verification of compiler correctness. The current reduction hypothesis is rejected. Automatic construction is implemented, but no measured advantage in modeling effort, semantic coverage or external tool cost has been established.

C.9 In-activity revision with digital and observed outcomes

A local embedding witness connects inspection, a compiled adapter-model preview and policy activation to an already-issued gate command. The accepted task, subject, deadline and command version persist across policy v1→v2. The pending command retains its accepted 50mm target, while new commands after the edit use 80mm; digital execution reaches that new target. The preview model assumes successful motion, so it cannot establish the physical outcome. Digital execution commits its own state directly. The device adapter receives explicit local report fixtures and never reads an unreported physical state. No robot, enterprise service or actual game plugin is involved.

Table S7: Delayed observation and physical-outcome evidence.

Evidence received at time 7; deadline 5	**Outcome**	**Meaning**
Command completed at 4	Timely effect, late confirmation	Task fulfillment and notification latency remain separate
Command completed at 5	Timely effect, late confirmation	Deadline is inclusive
Command completion event at 6	Late completion	Not counted as successful fulfillment
Only a positive state sample at 6	Unknown deadline status	The sample does not reveal first completion time
Acceptance with no completion evidence	Unknown	Acknowledgment is not success

An independent ordinary workflow adjudicator given identical reports returns identical statuses. Wrong task, subject, environment, role, command version, impossible timestamps and evidence type are rejected; a duplicate does not produce a second effect. A stale preview and a physical rollback request are rejected. The complete local run takes about 0.016 seconds wall time and 0.014 seconds CPU, including compilation and binding. The earlier 0.047-second interface fixture retained only a policy label and used a garden preview unrelated to that edit; it is preserved separately and is not evidence of executable gate-policy change. All new experiments use zero external model/API calls; preparation cost is unknown. The result establishes two distinct authority modes and a usable local interface, not unique cross-domain capability, entry-point benefit or learned transfer.

C.10 Shared-source integration and delivery audit

The shared-source episode connects the previously separate local mechanisms. Two scripted subjects use the frozen HTTP/A2A subset to propose and confirm their obligations. The authoritative world emits a three-unit heat effect at tick 2, due at tick 6. A subsequent HTTP task carries a material rule edit. The same backend object supplies proposed-world preview and, on acceptance, actual execution; the two subjects make no further requests until returning at tick 12. The pending heat, original goals and source identities persist. An installed, predefined resonator receives the old heat and is subsequently used by a helper. This is executable creation of an available affordance through installation, not intelligent invention of an instrument.

Twelve fixed executions compare resumable compilation and direct interpretation with identical rule source, finite plans, cache opportunities and information. Each method fulfills 12 of 12 old obligations and completes two of six expansions. Two infeasible and two stale edits per method do not complete the new activity. All six paired state trajectories agree. The 4.114-second wall and 0.760-second CPU totals include the full runner; there are no external model calls. A supplemental delivery audit makes the required three-task client set explicit after finding that the original audit assumed a complete supplied list. Original data and scores remain preserved; the strengthened check accepts the actual twelve deliveries and rejects missing clients, wrong bindings, all-rejected clients and missing wire responses.

This closes a local implementation gap, but does not establish that shared semantics has a unique modeling benefit: the equally capable shared-source interpreter needs no separate per-edit dynamics mapping either. It is a single-authority world, not a combination of live rule migration with the earlier partition-crash experiment or the physical-report fixture.

C.11 Native-host integration, development failures and costs

The fixed protocol, complete source snapshots, adapted native functions and all execution traces are at `next-research/2026-11-21-host-adaptation/fixed-01/`; the preceding shared-source episode is fixed at `next-research/2026-11-21-shared-source/fixed-01/`.

The next study executes the actual earlier handwritten garden transition through a new explicit adapter. The native host keeps its own control flow, nested promise records, interaction logic and internal event order. The adapter generates two native scalar edits and observation hooks, and supplies the state/event correspondences described in main Section 4.3. It is allowed the same source, checks, reuse and information as the compiler condition. It is not required to handwrite a fresh conversion for each case. Both conditions are restricted to the stated two-control edit domain.

The complete task now changes two interacting parts of the activity: the resin ignition threshold from 2 to 4 and helper use from tick 7 to tick 9. Heat is already in flight; old escort and habitat obligations remain. A successful receipt requires actual installation, heat reception and use at the requested time. Four conditions—safe two-part edit, unsafe threshold-3 boundary, a 0.20-second injected analysis delay and an injected old-helper-time mapping error—cross two winds and two implementations, for sixteen predeclared development executions. The mapping fault is applied equally to both backends and is a checking control, not an observed human mistake or an intentionally weakened baseline.

Table S8: Native-host activity outcomes.

Outcome per method, eight executions	**Declaration compiler**	**Native-host integration**
Old obligations fulfilled / required	16 / 16	16 / 16
New activities completed / requested	2 / 8	2 / 8
Unsafe edits refused	2	2
Stale edits refused	2	2
Injected mapping faults refused	2	2

All successful new activities use the instrument at tick 9. The safe compiler executions activate at tick 2 and install at tick 3; the safe native executions activate at tick 3 and install at tick 4. Consequently those paired state histories differ. The other six pairs have equal state histories, and all eight pairs have equal final activity outcomes. Late decisions occur at ticks 10–12 and are retained as observed. Refusing an expansion does not satisfy that expansion; old fulfillment is reported separately.

The fixed runner takes 5.064 seconds wall time and 0.801 seconds CPU. Summed construction, conformance-check, preview and actual-step times are respectively 0.026, 0.056, 0.019 and 0.009 seconds for the compiler, and 0.090, 0.109, 0.056 and 0.021 seconds for the instrumented native host. These overlapping internal components do not sum to wall time. Observation hooks add native overhead, and the short timings are not a scalability or production-engine ranking. Development runs add 0.794 seconds wall time, including one preserved native-construction failure caused by an AST field-name typo; scorer controls add 0.023 seconds. Authoring and preparation effort remain unmeasured, not zero.

After initial integration, both implementations consume the same two-value source edit without another handwritten per-case dynamics mapping. However, the native rule anchors, nested promise projection, version fields and event correspondence were written for this host. The unchanged declaration hash did not prevent an injected tick-7 execution from violating the tick-9 request; the shared bounded checker caught it. This supports a precise adapter trust boundary and a working full episode, but neither an independent reduction in per-edit labor nor a higher complete-activity success rate. It does not measure arbitrary host transfer, external benchmark performance, human productivity, learned mechanisms, real devices or partitioned hot updates.

D Extended prior-work analysis

Generative Agents connects memory, reflection, and planning to believable social behavior [1]. Concordia grounds natural-language agent actions through a game master that maintains state, generates observations, and resolves consequences [2]. Thus, separating agent proposals from world outcomes is established practice. AgentSociety explicitly studies independent asynchronous agents and scalable societal environments [3]. Emergence World studies persistent heterogeneous populations and governance, with subsequent work examining failures that propagate through memory and interaction [4, 5]. Long duration and multi-agent scale alone do not distinguish Fuxi.

SpatialOS provides a historical distributed-world reference [6]. Project Eden separates evolving structured state, state-conditioned observations, and generative rendering [7]. Its public description is a research preview, not a complete validation of every runtime property. FlockWorld studies synchronized multi-agent video generation [8]; visual consistency and social or resource consistency require different evaluations.

POET co-evolves environments and solvers [9]. Voyager studies skill accumulation [10], while SIMA2 and Dreamer4 connect agents, virtual environments, and world-model-based learning [11, 12]. These works rule out broad novelty claims for environment generation, skill reuse, or learning inside a model. Formal specification and invariant-confluence research provide established tools for reasoning about execution and coordination [13, 14].

Affordance modeling is an especially close neighbor. AGWM tracks action prerequisites through dynamic graph states; its published formulation uses a fixed environment-specific DAG schema and notes partial observability as future work [15]. Partial world modeling conditioned on affordances also has a formal foundation [16]. Fuxi must therefore distinguish changing the executable semantics of a running world from discovering capabilities inside a fixed schema. Virtual Institutions already connects human and agent participation with regulated virtual environments [17]. Project Sid investigates collective rules and cultural transmission [18], and RoboGen generates environments and supervision for skill learning [19]. These comparisons place the research question at persistent semantic extension and subsequent situated reuse, with equal extension authority required in the comparison.

Programmable World Model (PWM) is a direct recent comparator: executable entity rules maintain explicit persistent state, while state-augmented 3D boxes condition a generative renderer [22]. Its method supports modifying behavior through instructions, and its executor explicitly allows changes to the rule set. We therefore cannot claim separation of state and rendering, programmable interactions, or editability as new. Its reported benchmark emphasizes state–video consistency; a comparison with Fuxi would need to add compositional intervention, situated prediction, and persistent-extension tasks without restricting PWM's programmable engine.

Object-centric and relational prediction also has strong precedents. Interaction Networks, Graph Network-based Simulators, and MeshGraphNets learn physical dynamics through structured interactions [23, 24, 25]. Compositional message passing or generalization to more objects cannot alone distinguish Fuxi. A convincing learned extension must compete with these inductive biases and a precise classical simulator wherever the rules are known.

Dynamic software updating is another direct neighbor. Kitsune explicitly supports update points and migration of program state and control [20]. INVEST reuses verification results after component changes [21]. Therefore, state migration, live code replacement, and dependency-based revalidation are inherited methods. A strong extensible runtime can already implement them. Fuxi's possible research contribution must be assessed at the level of agent-visible evolving semantics, continuity contracts, and situated adoption, without claiming a new update or caching algorithm.

Self-modifying execution systems provide further direct comparators. Meta-Harness searches harness code using prior candidates, scores, and raw execution traces [28]. Its text-classification ablation favors trace access over scores or summaries, but that finding is task- and proposer-dependent. Its coding experiment searches and evaluates on the same 89-task benchmark, so those results do not establish independently held-out task transfer. DGM already evolves executable agent variants and maintains a candidate archive [29]. Hyperagents makes both the task agent and meta-level modification procedure editable [30]. Editing a runtime, preserving traces, or allowing the updater to change is therefore not a Fuxi-specific innovation.

Adaptive Auto-Harness separates persistent evolution workspaces from deployable harnesses [31]. That separation informs our design, but its v3 Appendix B.1 reports outcome-bearing trajectory-index labels and unresolved premature-read effects in retrospective prediction records. We do not use its reported gains as evidence of prospectively isolated forecasting. None of these external results has been reproduced here.

The local research program also contains negative evidence. Earlier event-prediction experiments did not establish a stable advantage over strong classical models. A shared action/editing representation evaluated in two grid engines did not achieve a joint performance advantage sufficient for promotion. Those experiments did not evaluate open-ended creation or social continuity. We inherit their constraints without treating them as results of this paper: shared representations are hypotheses, and reliable classical planning remains a mandatory baseline.

SpatialOS's official archived RPG example uses Unreal as a worker, and its experimental database-sync worker stores cross-session game data [33]. These artifacts support a historical architectural comparison, not a claim about current service availability or a reproduced scale result. Orleans already addresses virtual actor placement, lifecycle and recovery [34]. Reliable actor systems with application-level causal IDs, transactions, durable messages and resource owners are therefore necessary strong controls. A weak fire-and-forget implementation cannot establish Fuxi's novelty.

A2A v1.0.1 provides the interoperability reference [32]; our additional world profile is explicitly application semantics. Distributed simulation, logical causality and recoverable entity/message workflows have established foundations [35, 36, 37]. Growing Neural Cellular Automata shows a different form of local learned growth and regeneration [38]; it does not establish the ecology or resource consistency of a persistent shared world. Combining these precedents is a design proposal whose independent scientific value remains to be measured.

Live synthesis already addresses residual obligations during updates [39]; GR(1) update research supplies transition-controller mechanisms [40]. Compositional on-the-fly LTLf synthesis performs

game solving and reduction together [41], and hierarchical reactive synthesis uses local games under explicit locality conditions [42]. We inspected the relevant definitions and method portions, not complete external replications. These works rule out claiming obligation handover, bridge planning or compositional reduction in general as new.

Automatic update-model construction also predates this work. Hayden et al. merge an old program and its patch into a program suitable for existing verifiers and state an equivalence result for their formal setting [43]. Timing-aware PLC model generation provides another established compilation precedent [44]. Our finite compiler is therefore an implementation contribution to this world-runtime prototype, not a first semantic compilation method. Its comparison grants a generic compiler the same source, time information, constant propagation, common-subexpression cache and query roots.

E The author's AI1.0–4.0 research frame

The broader research program uses AI1.0–4.0 as an author-defined conceptual and science-fiction-motivated frame. It is not a recognized historical periodization, an established capability scale, or a claim that the present system implements its later stages. The frame motivates research questions; empirical claims in this paper use operational definitions independent of its numbering.

AI1.0: present work. The originating account describes what is currently being done without specifying an exhaustive technical definition. We therefore do not equate this label with statistical learning, language models, or any other exclusive class of methods.

AI1.5: world modeling. Fuxi supplies a setting for predicting how objects, agents, and events evolve under a declared world version and available observations. Predictive accuracy, uncertainty calibration, long-horizon consequences, and decision usefulness are measurable. A world model can already incorporate causal structure; this label does not imply that all world modeling is merely correlational.

AI2.0: causal modeling. The additional demand is explicit meaning and evidence for interventions, mechanisms, and counterfactual queries. A learned prediction conditional on an observed action need not equal the effect of assigning that action. Conversely, a known digital rule can directly define an intervention in its synthetic domain. Fuxi must declare which case applies and what stays invariant when an intervention or mechanism change is proposed. main Section 2.5 makes the distinctions precise.

AI3.0: shared human–agent life changing exploration. Humanistic concerns, emotion, impulse, and creative intuition may bring an unanticipated activity, objective, or possibility into consideration. Fuxi provides a persistent environment where participants can enact, modify, and revisit such proposals. For example, a rain-interrupted cooking task might lead participants to create a communal music activity; it must not silently count as completing the original cooking task. Nuwa interprets and acts, Pangu supports creation, and Taihe examines the value of the changed direction. This is a research possibility, not evidence that irrationality is generally beneficial or that noise generates valuable novelty.

AI4.0: AI designing AI. The author's scope includes mathematical foundations, technical structures, and potentially physical materials. Fuxi can supply versioned test worlds, controlled interventions, and records for evaluating proposed improvements to predictors, tools, or improvement procedures. A runtime that hosts these tests is not itself AI4.0, and changing a world rule is not automatically improving an AI. Designing a better architecture, changing an improvement procedure, and recursively using that procedure are distinct claims. An AI-designed system requires evidence for its specific redesign and outcomes. The narrower claim of a better improvement procedure requires downstream gains beyond additional search or compute; recursive use additionally requires evidence about sustained effects and costs across iterations. The present experiments establish none of these claims. The frame does not require AI3.0 to be a necessary historical predecessor of AI4.0.

The positive objective across these questions is a world in which people and agents can play, cooperate, create, remember, and discover worthwhile new activities. The runtime's reliability mechanisms preserve the consequences of that life; they are not its sole purpose.

F Broader discriminating evaluation design

A deployment study must go beyond local parallel service: multiple physical hosts, replicated transport and directory failure, real bandwidth and disk limits, prolonged operation, and sustained offered-load sweeps. Report object and active-agent counts, cascade depth and fan-out, accepted versus rejected work, lost and duplicated effects, latency tails, recovery time, CPU, memory, storage and retained history. Partition, retry, cancellation, re-entry and semantic upgrades must be exercised together. Compare a capable actor/runtime and direct-task baseline under equal agent policies and semantics. Protocol conformance, shared-world correctness and cooperative benefit remain separate outcomes. Unimplemented features in the implementation matrix are evidence gaps, not zero-cost successes.

A complete dynamics study first separates known mechanics, hidden parameters, unseen object counts/topologies, held-out mechanism compositions, and genuinely new semantic operators. Compare exact classical simulation where available, learned graph simulators, a programmable explicit-state world such as PWM with equivalent rule authority, and version-conditioned causal prediction. Intervention targets, initial states, and uncertainty budgets must match. Measure multi-step prediction, intervention-response accuracy, failure calibration, outcome controllability, discovery under limited observation, and computational cost. Do not define the test simulator using one model's assumptions and call that model's agreement an independent benchmark.

The decisive learning hypothesis is that structured, version-aware consequence prediction improves competent action after unfamiliar combinations or mechanism changes. It is falsified if a symbolic registry plus planner or a standard relational predictor matches it under equal observations and budget. The current hand-authored diagnostics do not test this hypothesis.

A complete study should also include entering a world, cooperating, creating a mechanism, leaving, returning, and transferring a relevant capability to a semantically different world. Navigation and resource tasks are components rather than substitutes for this sequence.

The decisive runtime comparator must support live semantic extensions, migration, contract checking, and standard incremental verification. Removing those capabilities from the baseline would predetermine the outcome. Baselines additionally include strong classical execution and planning, a frozen foundation agent with the same tools and memory, learned event prediction, dynamic-affordance modeling, semantic extension, and finally coupled world/agent curriculum adaptation. Every condition must have equivalent creative authority. Giving one condition fixed actions and another arbitrary extensions would confound an architectural intervention with unequal capability access.

Splits should separate mechanism families, partner-policy lineages, author templates, and environment semantics. Shared representations must be compared with independently parameterized models trained on the same data. Editing gains cannot be assigned to the action policy when a classical planner actually executed the edited world. Failed calls, interventions, and training costs remain in the accounting.

Primary outcomes concern valid realized cooperation, adoption and reuse of new mechanisms, preservation violations, transfer, and cost. Human studies separately examine whether participants value their choices, recognize continuity, and want to continue shared activities. Return rate alone is not a measure of care or autonomy. World-level interference requires cluster assignment and uncertainty estimates at the world or social-group level; action counts are not independent sample sizes.

A component comparison should distinguish the runtime's contribution from generation and learning. Hold generated mechanisms, agent policies, observation budgets, and replayable exogenous events fixed when testing runtime changes. Hold the runtime fixed when testing migration-aware prediction or agent adaptation. Discovery and adoption require fresh partners without access to creator-private traces. A strong symbolic registry and planner must receive the same public mechanism descriptions as a learned system. Shortcuts that bypass resource costs are not creative improvements unless such changes are explicitly within the permitted design space.

For the author's AI3.0 question, evaluation must retain the original activity, the proposed new direction, who proposed it, what participants chose, and subsequent voluntary use. Open-ended value inquiry can admit new goals without retroactively declaring failure on the original goal a success.

Compare shared discovery with fixed-goal optimization and with an equally capable co-creative baseline; distinguish genuine participant endorsement from model-generated approval. No such participant study has been run here.

For the AI4.0-motivated question, compare fixed and editable improvement procedures from the same initial task system, under equal total resources on untouched problem families. Report multi-round learning curves, old-task retention, cross-task interference, deployment-selection failures, and the complete costs of data preparation, updater calls, retrieval, validation, failed attempts, and serving. Keep the original evaluator independent. This tests a bounded hypothesis about an improvement process; it cannot establish the entire speculative stage.

The proposal loses support if a simpler system with matched tools and memory performs equally well, if new options are not adopted, or if short-term gains destroy prior relationships or transfer. An honest negative result should revise the architecture rather than change the evaluation target after seeing outcomes.

References

[1] Park et al. Generative Agents. https://arxiv.org/abs/2304.03442

[2] Vezhnevets et al. Concordia. https://arxiv.org/abs/2312.03664

[3] Piao et al. AgentSociety. https://arxiv.org/abs/2502.08691

[4] Akkil et al. Emergence World: A Platform for Evaluating Long-Horizon Multi-Agent Autonomy. https://arxiv.org/abs/2606.08367

[5] Akkil et al. Emergence World: Adversarial Stress-Testing of Long-Horizon Multi-Agent Systems. https://arxiv.org/abs/2609.17320

[6] Improbable. SpatialOS UnrealSDK introduction. https://github.com/spatialos/UnrealSDK/blob/master/docs/introduction.md

[7] VAST AI Research. Project Eden. https://www.tripo3d.ai/blog/project-eden

[8] Koksal et al. FlockWorld. https://arxiv.org/abs/2610.34564

[9] Wang et al. POET. https://arxiv.org/abs/1901.01753

[10] Wang et al. Voyager. https://arxiv.org/abs/2305.16291

[11] SIMA Team. SIMA2. https://arxiv.org/abs/2512.04797

[12] Hafner et al. Training Agents Inside of Scalable World Models. https://arxiv.org/abs/2509.24527

[13] Lamport. The Temporal Logic of Actions. https://www.microsoft.com/en-us/research/publication/the-temporal-logic-of-actions/

[14] Bailis et al. Coordination Avoidance in Database Systems. https://arxiv.org/abs/1402.2237

[15] Zhang et al. AGWM. https://arxiv.org/abs/2605.06841

[16] Khetarpal et al. Affordances Enable Partial World Modeling with LLMs. https://arxiv.org/abs/2602.10390

[17] Bogdanovych, Simoff, and Esteva. Virtual Institutions Prototype. https://www.ifaamas.org/Proceedings/aamas09/pdf/06_Demos/d_02.pdf

[18] Altera.AL et al. Project Sid. https://arxiv.org/abs/2411.00114

[19] Wang et al. RoboGen. https://arxiv.org/abs/2311.01455

[20] Hayden et al. Kitsune: Efficient, General-purpose Dynamic Software Updating for C. OOPSLA 2012. https://www.cs.umd.edu/~mwh/papers/kitsune.pdf

[21] Johnson, Calinescu, and Kikuchi. An Incremental Verification Framework for Component-Based Software Systems. CBSE 2013. https://eprints.whiterose.ac.uk/id/eprint/75438/1/cbse21_inc.pdf

[22] Huang et al. Programmable World Model. arXiv:2609.10540v2, 2026. https://arxiv.org/abs/2609.10540

[23] Battaglia et al. Interaction Networks for Learning about Objects, Relations and Physics. 2016. https://arxiv.org/abs/1612.00222

[24] Sanchez-Gonzalez et al. Learning to Simulate Complex Physics with Graph Networks. 2020. https://arxiv.org/abs/2002.09405

[25] Pfaff et al. Learning Mesh-Based Simulation with Graph Networks. 2020. https://arxiv.org/abs/2010.03409

[26] Pearl. An Introduction to Causal Inference. The International Journal of Biostatistics, 2010. https://ftp.cs.ucla.edu/pub/stat_ser/r354-reprint-corrected.pdf

[27] Shpitser and Pearl. Complete Identification Methods for the Causal Hierarchy. JMLR 9, 1941–1979, 2008. https://www.jmlr.org/papers/v9/shpitser08a.html

[28] Lee et al. Meta-Harness: End-to-End Optimization of Model Harnesses. arXiv:2603.28052v1, 2026. https://arxiv.org/html/2603.28052v1

[29] Zhang et al. Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents. arXiv:2505.22954v3. https://arxiv.org/abs/2505.22954v3

[30] Zhang et al. Hyperagents. arXiv:2603.19461v2, 2026. https://arxiv.org/abs/2603.19461v2

[31] Liu et al. Adaptive Auto-Harness: Sustained Self-Improvement for Agentic System Deployment on Open-Ended Task Streams. arXiv:2606.01770v3, 2026. https://arxiv.org/html/2606.01770v3

[32] A2A Project. Agent2Agent Protocol Specification v1.0.1. Accessed 20 November 2026. https://a2a-protocol.org/v1.0.1/specification/

[33] Improbable. SpatialOS RPGDemo and Database Sync Worker (archived official examples). https://github.com/spatialos/RPGDemo ; https://github.com/spatialos/database-sync-worker

[34] Bernstein et al. Orleans: Distributed Virtual Actors for Programmability and Scalability. MSR-TR-2014-41. https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/Orleans-MSR-TR-2014-41.pdf

[35] Chandy and Misra. Distributed Simulation: A Case Study in Design and Verification of Distributed Programs. IEEE TSE, 1979. https://www.cs.utexas.edu/~misra/scannedPdf.dir/DistrSimulation.pdf

[36] Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. CACM, 1978. https://lamport.azurewebsites.net/pubs/time-clocks.pdf

[37] Helland. Life beyond Distributed Transactions: an Apostate's Opinion. CIDR, 2007. https://www.cidrdb.org/cidr2007/papers/cidr07p15.pdf

[38] Mordvintsev et al. Growing Neural Cellular Automata. Distill, 2020. https://distill.pub/2020/growing-ca/

[39] Finkbeiner, Klein, and Metzger. Live synthesis. Innovations in Systems and Software Engineering 18, 443–454, 2022. https://link.springer.com/article/10.1007/s11334-022-00447-5

[40] Dynamic Update for Synthesized GR(1) Controllers. ICSE 2022. https://smlab.cs.tau.ac.il/syntech/dupdate/dupdate-icse22-cr.pdf

[41] A Compositional Framework for On-the-Fly LTLf Synthesis. arXiv:2508.04116v2. https://arxiv.org/html/2508.04116v2

[42] Dynamic Hierarchical Reactive Controller Synthesis. arXiv:1510.07246v2. https://arxiv.org/pdf/1510.07246v2

[43] Hayden et al. Specifying and Verifying the Correctness of Dynamic Software Updates. Author preprint. https://www.cs.tufts.edu/~jfoster/papers/vstte12.pdf

[44] Fernandez Adiego et al. Modelling and Formal Verification of Timing Aspects in Large PLC Programs. IFAC World Congress, 2014. https://cds.cern.ch/record/1956687/files/CERN-ACC-2014-0226.pdf?version=1

[45] Fu and Topcu. Integrating active sensing into reactive synthesis with temporal logic constraints under partial observations. arXiv:1410.0083. https://arxiv.org/pdf/1410.0083

[46] Capretto et al. Offchain Runtime Verification (for The Tezos Blockchain). Author manuscript; accessed 21 November 2026. https://software.imdea.org/~cesar/papers/capretto25offchain.pdf

[47] Microsoft. Compatible grains. Official Orleans documentation; accessed 21 November 2026. https://learn.microsoft.com/en-us/dotnet/orleans/grains/grain-versioning/compatible-grains

[48] Temporal. Roll out and pin Workflows. Official documentation source; accessed 21 November 2026. https://github.com/temporalio/documentation/blob/main/docs/production-deployment/worker-deployments/worker-versioning/roll-out-and-pin.mdx

[49] Fu, Dimitrova, and Topcu. Abstractions and sensor design in partial-information, reactive controller synthesis. ACC 2014. https://eprints.whiterose.ac.uk/id/eprint/156437/1/acc14.pdf

[50] Microsoft. Event Sourcing pattern. Azure Architecture Center; accessed 21 November 2026. https://learn.microsoft.com/en-us/azure/architecture/patterns/event-sourcing