Human in Master Control

MCE-1: a minimum complete synthetic economy with constitutional causality and non-sovereign AI

Hayden Lindley · Glyphd Labs

2026-07-27

Contents

Human in Master Control

Abstract

Economic and institutional systems are often modeled as markets, optimization problems, or agent simulations. Those models can generate interesting behavior while leaving critical questions undefined: who has standing, which rights cannot be traded, how authority is delegated, what happens when evidence conflicts, how appeals work, when emergency power

expires, and whether an optimization engine can rewrite the conditions of human participation.

MCE-1 — Minimum Complete Economy is a software-only simulation contract that begins with institutional completeness rather than predictive realism. It models people and organizations, rights and needs, resources and ecological limits, capabilities, projects, work, contribution, claims, obligations, exchange, bounded capital, commons, automation gains, governance, adjudication, concentration, uncertainty, external institutions, branches, failure, and recovery.

The system may begin with synthetic people, random demand, simplistic equations, fictional institutions, and fake currencies. It may not begin with missing rights, missing authority, missing appeals, missing receipts, hidden callbacks, or a privileged sovereign operator.

Language models interpret, explain, propose, and simulate. They do not become the unappealable judge or canonical state authority. LCSM governs recognized state. FOVEA addresses and retrieves it. `fold/2` represents it. SurfaceMind presents it. PolyView synchronizes projections. Gantry governs crossings. A constitutional layer constrains every operation.

1. Why “minimum complete”

A simulation can be large and still structurally incomplete.

It may include:

- firms;
- prices;
- workers;
- demand;
- production;
- money.

Yet omit:

- human rights;
- sufficiency;
- disputes;
- jurisdiction;
- appeal;
- ecological limits;
- concentration;
- evidence;
- uncertainty;
- emergency expiration;
- external law;
- recovery.

The missing parts are often treated as later policy. MCE-1 treats them as system primitives.

“Minimum” means the simplest implementation that includes the whole institutional causal loop. “Complete” does not mean realistic prediction of Earth.

2. Terminal condition

A conforming world permits people, communities, organizations, and explicitly delegated software agents to:

- declare needs and proposals;

- discover resources and capabilities;
- form projects;
- perform and verify work;
- compensate contributors;
- finance activity without permanent capital sovereignty;
- exchange ordinary goods and services;
- allocate essentials by rights and need;
- distribute automation gains;
- replenish commons;
- govern affected systems;
- dispute facts and decisions;
- appeal;
- fork institutions under declared rules;
- preserve lineage and uncertainty;
- compare state with constitutional and ecological constraints.

Every operation remains replayable and attributable.

3. Human standing

The first constitutional invariant is unconditional human standing.

A recognized person does not lose standing because of:

- wealth;
- productivity;
- employment;
- reputation;
- debt;
- citizenship category;
- model score;
- political agreement.

The sufficiency floor cannot be revoked merely because a person is economically weak or unpopular.

This separates economic output from human worth.

4. Separation of powers and scores

The system maintains separate objects for:

- purchasing power;
- rights;
- contribution evidence;
- stewardship;
- obligations;
- authority;
- reputation/confidence;
- governance participation.

No scalar may control all of them.

This prevents wealth, popularity, or a model-generated score from silently becoming political and moral sovereignty.

5. Bounded capital

A capital instrument declares:

```
principal
risk
repayment priority
maximum return
governance rights
duration
loss allocation
conversion
transfer restrictions
termination
```

Capital cannot silently convert into perpetual control.

This is not a claim that one financial design is universally correct. It is a requirement that power created by capital be explicit, bounded, and terminable.

6. Contribution and lineage

Useful contribution remains addressable across:

- fork;
- merge;
- reuse;
- maintenance;
- correction;
- synthesis;
- derivation.

Lineage compensation considers materiality, directness, substitutability, persistence, evidence quality, distance, and contribution type.

It must also avoid infinite royalty chains through decay, caps, aggregation, expiration, and challenge.

7. Canonical state

MCE-1 distinguishes:

- world state;
- recognized state;
- believed state;
- observed state;
- disputed state;
- simulated state;
- external state;
- canonical state.

An observation does not become canonical because it is repeated, asserted by a powerful actor, assigned high model confidence, or useful to an optimal plan.

Truth promotion follows declared validation.

8. Universal object and event envelopes

A canonical object includes identity, schema, address, branch, jurisdiction, authority, confidence, visibility, lineage, dependencies, evidence, policy, events, supersession, and integrity.

A consequential event includes source, target, time, branch, jurisdiction, authority, priority, confidence, evidence, causal parents, idempotency, privacy, expected effects, and failure behavior.

The governed lifecycle is:

```
OBSERVED
→ INTERPRETED
→ PROPOSED
→ SIMULATED
→ AUTHORIZED
→ COMMITTED
→ EXECUTING
→ VERIFIED
→ SETTLED
→ APPEALABLE
→ FINALIZED
```

Divergence includes rejected, failed, disputed, rolled back, superseded, quarantined, and unresolved.

9. Economic RTOS

MCE-1 behaves as an economic real-time operating system.

It provides:

- clocks;
- priority lanes;
- queues;
- interrupts;
- deadlines;
- reservations;
- conflict control;
- process isolation;
- watchdogs;
- escalation;
- degraded modes;
- receipts;
- replay.

Priority begins with immediate life safety, rights/sufficiency, ecological boundaries, and critical infrastructure before ordinary exchange and optimization.

Lower lanes cannot permanently starve, except under bounded declared emergency.

10. Cells

A Cell is a recursive economic and governance domain:

- person;
- household;
- project;

- organization;
- neighborhood;
- city;
- watershed;
- market;
- nation;
- federation;
- ecological system.

Cells can overlap. A water plant may belong to municipal, watershed, infrastructure, cooperative, and emergency contexts without duplicating the object.

Each Cell declares rights, needs, resources, capabilities, projects, claims, obligations, governance, external relations, metrics, interrupts, and integrations.

11. Engines and separation

The architecture contains engines for:

- needs;
- sufficiency;
- resources;
- capacity;
- projects/work;
- allocation;
- exchange;
- settlement;
- governance;
- adjudication;
- concentration;
- resilience;
- evidence;
- external bridges;
- alignment control.

Money does not become the only signal. A language model does not become the final controller.

12. LLM role

Natural language compiles into typed Glyph Notes:

```
speaker
intent
target
scope
jurisdiction
requested_action
constraints
authority_claim
evidence
context
privacy
approvals
operations
unresolved_questions
```

The model can:

- interpret;
- synthesize evidence;
- explain;
- create scenarios;
- generate proposals;
- prepare typed actions.

It cannot:

- create its own standing;
- expand its authority;
- alter rights directly;
- invent evidence;
- conceal its principal;
- finalize unappealable adjudication;
- convert delegation into ownership.

13. Gantry and external systems

External systems—banks, sensors, logistics, law, model providers—remain outside until represented through a governed bridge.

Gantry enforces:

- schema;
- least authority;
- source identity;
- privacy;
- rate/capacity;
- failure;
- retry;
- receipts;
- revocation;
- quarantine.

No hidden callback may alter canonical state.

14. Shadow worlds

Candidate policy can run in a branch:

```
base branch
changed assumptions
changed policy
model versions
random seeds/uncertainty
time horizon
evaluation
```

High-impact changes can move from sandbox to shadow, limited Cell, bounded pilot, monitored expansion, or rollback.

The controller may adjust only parameters explicitly delegated to it. It cannot repeal rights or expand itself.

15. Water crisis trace

The reference vertical begins with conflicting water observations.

The system:

1. records observations and uncertainty;
2. detects contradiction;
3. estimates threatened needs;
4. raises an essential-shortage interrupt;
5. exposes ecological limits;
6. enumerates capacity;
7. proposes temporary allocation;
8. generates human-readable plans and typed operations;
9. validates external actions;
10. runs candidates in shadow worlds;
11. rejects constitutional or ecological violations;
12. obtains required authorization;
13. creates projects and work;
14. reserves resources;
15. executes;
16. verifies work;
17. allocates essential water by rights/need;
18. settles compensation and bounded capital;
19. handles a contribution dispute with interim continuity;
20. permits appeal;
21. compares predicted and actual outcomes;
22. modifies or rolls back policy;
23. expires emergency authority;
24. records a postmortem;
25. preserves replay.

The trace is complete only if no undefined institutional gap remains.

16. Mission Control

A reference surface contains:

- zoomable world/cell canvas;
- lenses for needs, rights, resources, production, exchange, work, ownership, lineage, commons, ecology, governance, disputes, concentration, risk, evidence;
- constitutional status strip;
- object inspector;
- branch/timeline ribbon;
- policy laboratory.

Every displayed state labels whether it is synthetic, inferred, disputed, external, counterfactual, or canonical.

17. Acceptance gates

Required gates include:

- constitutional completeness;
- ontological completeness;
- transition completeness;
- rights and authority;

- full water-crisis vertical;
- failure/recovery;
- deterministic replay;
- privacy/integrity;
- model non-sovereignty;
- observability.

Synthetic simplicity is allowed. Missing institutional causality is not.

18. Evidence state

The MCE-1 system contract is a large normative specification with object ontology, event protocol, RTOS, governance, plugs, tunnels, hooks, privacy, LLM contracts, UI, end-to-end trace, failure/recovery, and acceptance harnesses.

It is `SPECIFIED`, not evidence that the economy has been implemented or that its policies are empirically optimal.

19. Security and capture

Threats include:

- privileged founder/operator exemptions;
- model sovereignty;
- data poisoning;
- governance capture;
- wealth concentration;
- emergency permanence;
- hidden external bridges;
- identity and reputation collapse;
- privacy abuse;
- replay of authority;
- simulation presented as prediction.

The constitution and audit system must apply to operators as well as ordinary participants.

20. Limitations

Institutional completeness does not guarantee justice, stability, or realism. Values conflict. Rights and sufficiency definitions require legitimate human governance. Simulation agents can encode bias. Economic models can produce false confidence. Real legal systems may conflict with the synthetic contract. Complexity can obscure accountability.

MCE-1 is a research environment for making institutional assumptions executable and testable—not a unilateral replacement for existing society.

21. Conclusion

The central design choice is to make institutional causality explicit before making the simulation realistic.

MCE-1 can begin with fake money and random demand. It cannot begin with fake rights, hidden authority, or unappealable software power.

The final contract is:

An economy can be software-operable without making software sovereign over human standing.

Architecture illustration briefs

- **P21-F01 — Constitutional economy stack**: Constitution constrains LCSM, FOVEA, OOC, SurfaceMind, PolyView, Zeke, Gantry, RTOS, engines, and external tunnels.
- **P21-F02 — Governed transition lifecycle**: Observed through Finalized with divergent failed/disputed/rollback states and receipts.
- **P21-F03 — Water crisis vertical**: Forty-step reference trace condensed into observation, need, policy simulation, authorization, work, allocation, dispute, appeal, control feedback, and postmortem.
- **P21-F04 — Separation of scores and powers**: Rights, money, reputation, contribution, stewardship, authority, and governance remain separate objects.

Source register

- **P21-S01** — MCE-1 System Contract v0.1 (2026-07-25). `MCE-1-SYSTEM-CONTRACT-v0.1.md`
- **P21-S02** — GlyphCAS Design Report (2026-07-24). `How_we_can_make_CAS_lightning_fast_for_any_LL`
- **P21-S03** — Glyphd V1 authority, aliveness, and threshold consent invariants (2026-07-18). `GLYPHD-V1-MASTER-SPEC-LOCKS-AND-EXPERIMENT-PLAN-v0.1.md`
- **P21-S04** — LCSM Workbench E2E Specification (2026-07-04). `LCSM-WORKBENCH-E2E-SPEC.v0.md`

Public-disclosure and IP boundary

This edition publishes the public-safe institutional architecture and constitutional principles. It excludes private economic scenarios, unpublished policy parameters, confidential integrations, personal data, and patent-claim drafting. It is not a prediction of real economies, a legal system, an investment instrument, or a claim of validated social outcomes.