

Title

A New Epoch of Modeling & Simulation: Constitutional Governance and MVC² for LLM-Spawnd Agent-Based Simulation

Authors and Affiliations

Charles F. Day — Charles F. Day & Associates LLC (CFDay) / QWISE LLC

Projected Track

Emerging Capabilities

Abstract (490 words)

The emergence of general-purpose large language models has pushed modeling and simulation into a new epoch. For the first time at scale, non-specialists can assemble virtual SME teams by prompting a single LLM to produce multiple specialized agents, enabling rapid agent-based exploration of problems that historically required dedicated M&S platforms, long development cycles, and formal training in agent-based modeling. An expanding population of users is already conducting what are functionally agent-based simulations across the full gamut of professional, scientific, and operational tasks, whether or not they recognize them as such.

This represents a fundamental discontinuity. The M&S community has spent decades developing rigorous frameworks for validation, verification, accreditation, and governance of simulation systems. Those frameworks assumed institutional actors operating purpose-built platforms within defined accreditation boundaries. That assumption no longer holds. We have entered an early, ungoverned expansion in which capabilities are compounding faster than governance, repeatability, and trust mechanisms. The same gap extends to adjacent quantum-inspired simulation stacks, where agentic AI supports both the simulation operators and the simulation itself. Without a common measurement standard, there is no principled way to determine whether any agentic system, classical or quantum-inspired, is producing trustworthy outputs.

This paper argues that constitutional governance is the missing layer for agentic M&S and proposes Minimum Viable Coherence and Cohesion, designated MVC², as the minimum composite metric to determine whether a multi-agent system remains within governed bounds. MVC² integrates two complementary dimensions: coherence, the degree to which agent outputs remain internally consistent with constitutional constraints, evidence requirements, and scenario assumptions; and cohesion, the degree to which a multi-agent team remains coordinated toward a shared objective without role drift, hidden authority escalation, or contradictory actions. When MVC² falls below threshold, the system has entered governance decoherence and its outputs should not be trusted for consequential decisions. MVC² is intended as a testable pass-fail threshold suitable for incorporation into governance standards including ISO/IEC 42001 and the EU AI Act, providing a quantitative measurement layer that

current frameworks do not consistently require. In practice, MVC² serves as a gating threshold for escalation, scope constraint, or reversion to human-directed operation. A separate submission describes an implementation architecture and operational enforcement mechanisms.

Operational evidence from a five-year federal services deployment demonstrates that governed agentic systems can scale without sacrificing auditability: within a defined problem scope and agent configuration appropriate to the task, nine agents dynamically instantiated at runtime produced a 12-slide analytic deliverable in 3.2 seconds with full lifecycle audit trail; a quantum-inspired portfolio optimization achieved 98 of 100 at 1,000-asset scale in 8.88 seconds with a 10,000-iteration audit trail; and in compliance evidence generation, human-on-the-loop workflows processed approximately 200 research-intensive questions at approximately 3 minutes per question, reduced from a 45-minute baseline (approximately 93 percent reduction), achieving 110 controls assessed with all 320 assessment objectives satisfied and zero deficiencies in independent C3PAO external audit. These results illustrate that the new epoch of agentic M&S demands, and can support, principled governance measurement.

PAPER #2

MODSIM World 2026 | Autonomy

v2 CLEAN SUBMIT | May 22, 2026

Title

Society of Agents: A Constitutional Governance Architecture for LLM-Spawnd Multi-Agent Simulation

Authors and Affiliations

Charles F. Day — Charles F. Day & Associates LLC (CFDay) / QWISE LLC

Alok Chaturvedi (Senior Scientific Advisor) — Simulex / Advisory role

Projected Track

Autonomy

Abstract (455 words)

Recent work proposes Minimum Viable Coherence and Cohesion (MVC²) as a standardizable composite metric for determining whether LLM-spawnd multi-agent systems remain within governed bounds. This paper presents the architecture that operationalizes that metric: the Society of Agents (SoA), a constitutional governance framework for multi-agent M&S in which agents are dynamically instantiated under constitutional governance and operated under SoA constitutional constraints within the defined test framework.

SoA enforces governance structurally, not through monitoring after the fact. The system is deterministically constrained within defined governance bounds, specifying role authorities, admissible actions, evidence requirements, and escalation conditions across human-in-the-loop (HITL) and human-on-the-loop (HOTL) modes. A governance-conditioned resolution gate ensures that every agent output passes through constitutional validation before inclusion in system results, producing auditable, reproducible decision pathways. A commercial LLM, operating as an adversarial stakeholder, probes constraint violations and surfaces latent assumptions through governed convergence—an iterative validation cycle that measures error modes and directs correction under human oversight.

The architecture extends to quantum-inspired simulation stacks, where quantum-inspired structural analogies applied to governance and coordination, implemented in classical systems, represent convergence behavior consistent with quantum-inspired annealing models. In this mode, a quantum-inspired portfolio optimization is supported by three layers of agentic AI with HITL governance and optional HOTL oversight, providing transparency, explainability, and agentic controls across the optimization lifecycle.

SoA operationalizes MVC² as a governance trigger. When MVC² is satisfied, the system operates with increasing AI operational responsibility within bounded constitutional controls. When

MVC² falls below threshold, the architecture enforces automatic escalation: increased human oversight, constrained agent scope, or full reversion to human-directed operation. This makes MVC² not merely a reporting metric but an enforceable governance boundary with deterministic consequences.

Within a defined problem scope and agent configuration appropriate to the task, we report two validation events demonstrating governed convergence under constitutional constraints. First, an adversarial validation in which a commercial LLM generated 88 structured test problems across 10 mathematical fields. Adversarial LLM evaluation indicated approximately 60 percent system accuracy. Initial performance confirmed that assessment. Human-on-the-loop observation identified a systematic error pattern; human-in-the-loop directed the system to self-assess and generate its own corrective instructions, which were incorporated as a new governance iteration. Post-correction accuracy reached 97.7 percent, with facet-level governance quality improving from 94.3 to 98.8 percent. The 40-percentage-point improvement quantifies the value of human-governed constitutional AI—accuracy achieved through governance intervention, not increased model capability. Second, in compliance evidence generation, HOTL workflows processed approximately 200 research-intensive questions at approximately 3 minutes per question, reduced from a 45-minute baseline (approximately 93 percent reduction), achieving 110 controls assessed with all 320 assessment objectives satisfied and zero deficiencies in independent C3PAO external audit. These results demonstrate that constitutional governance architecture enables trustworthy acceleration in the new epoch of agentic M&S.