

Simulation as Civic Infrastructure

Preference Revelation and the End of Narrative Coordination

Flyxion

February 4, 2026

Abstract

Contemporary civic and epistemic institutions remain largely organized around narrative coordination. Education systems, democratic mechanisms, and feed-based social media presume that collective understanding and preference formation can be achieved through discourse, persuasion, and symbolic alignment within a shared but largely static background. This essay argues that such institutions are increasingly misaligned with the structure of the systems they attempt to govern. As economic, technological, and ecological dynamics become high-dimensional, nonlinear, and strongly path-dependent, narrative forms of coordination fail to expose the consequences of choices or to bind preference expression to irreversible commitment.

The essay situates this institutional mismatch within a broader family of scientific critiques that reject narrative primitives in favor of relational, scale-free, and constraint-driven structure. Parallels are drawn to Barbour's denial of fundamental time, Penrose's conformal cyclic cosmology, and entropic approaches to physical law in which geometry, scale, and temporal direction emerge from irreversible processes rather than serve as foundational coordinates. Across these frameworks, explanation advances by abandoning stories told against fixed backgrounds and replacing them with selection among admissible structures under global constraints.

Against this backdrop, the essay argues that contemporary production has already crossed an epistemic threshold. Factory simulation, digital twins, and counterfactual optimization increasingly function as the primary coordination substrate of industrial activity, enabling global feasibility to be evaluated prior to material commitment. This simulation-native mode of reasoning renders discourse-centric institutions progressively obsolete, insofar as they continue to mediate collective choice through symbolic preference expression rather than consequence-bearing exploration.

The essay therefore proposes shared civic simulation environments, analogous in strategic form to long-horizon 4X simulations, as a replacement for feed-based coordination and narrative-driven education. Within such environments, macrostructural alternatives are explored as trajectories under explicit physical, economic, and normative constraints, and collective preferences are revealed through sustained interaction with consequences rather than through rhetorical alignment. On this view, simulation becomes a civic infrastructure that reorganizes legitimacy and collective agency around irreversible trajectories, rendering futures intelligible not as narratives to be believed, but as structures to be inhabited, tested, and selectively sustained under conditions of complexity, path dependence, and irreversibility.

1 Introduction

Many of the most influential explanatory systems of the modern world are organized as narratives. They assume a background that remains fixed while events unfold upon it, and they treat the relevant primitives as beginnings, sequences, and outcomes. In cosmology, the background is spacetime and the narrative begins with a Big Bang followed by expansion. In political economy, the background is a market environment and the narrative begins with scarcity followed by growth. In education and public discourse, the background is a stable social world and the narrative proceeds through curricula, persuasion, and credentialed advancement. Such narratives can be pragmatically useful, but their explanatory power depends on a hidden assumption: that the future is sufficiently low-dimensional that it can be coordinated by stories.

A different tradition has always existed beneath the narrative surface, insisting that what is fundamental is not story but structure. In physics, this tradition appears wherever scale, time, or geometry are treated not as primitives but as emergent constraints. Barbour's work denies time as a fundamental parameter and replaces temporal succession with a relational ordering of configurations, where duration is a derivative measure of change (Barbour 1999). Penrose's conformal cyclic cosmology removes absolute scale at cosmological boundaries, arguing that when mass becomes irrelevant, metric size loses physical meaning and only conformal relations persist (Penrose 2010). In the Relativistic Scalar Vector Plenum (RSVP) framework, expansion and even spacetime are treated as secondary descriptions of a deeper substrate of irreversible entropy flow and constraint relaxation, such that temporal direction is identified with thermodynamic irreversibility rather than with a coordinate time.

This essay argues that the economy is undergoing an analogous transition. Production is increasingly organized not by narrative instruments such as plans, slogans, or even price signals alone, but by simulation infrastructures that explore counterfactual trajectories and constraint surfaces before material commitment occurs. Digital twins, factory simulators, supply-chain solvers, and design-to-manufacture pipelines instantiate a mode of reasoning that is explicitly relational, path-dependent, and consequence-facing. The dominance of this mode in production implies that discourse-centric institutions become progressively misaligned with the epistemic requirements of collective life. If the central coordination substrate of production is simulation-native, then education and social media cannot remain primarily narrative without creating a widening gap between symbolic preference expression and causal leverage.

The essay therefore advances a further claim. When societies must choose among macrostructures whose consequences unfold over long horizons and through nonlinear dynamics, preferences cannot be meaningfully expressed as slogans, votes, or engagement metrics alone. Preference revelation must occur through sustained interaction with structured counterfactual spaces. A shared civic simulation environment, analogous in strategic form to long-horizon 4X games, can function as a medium for exploring competing futures, exposing tradeoffs, and selecting trajectories that exhibit robustness under explicit constraints. In this proposal, simulation does not trivialize politics by turning it into entertainment; it replaces discourse with consequence-bearing exploration, reconstituting legitimacy and collective agency around irreversible trajectories rather than around attention.

2 From Fundamental Narratives to Relational Structure

Modern scientific worldviews have repeatedly been organized around narrative primitives: a beginning, a sequence of transformations, and an end state governed by laws that act upon a fixed background. In cosmology this takes the form of an initial singularity followed by expansion; in

economics it appears as growth trajectories driven by incentives; in education and media it manifests as linear curricula and discursive persuasion. These narrative forms are not merely stylistic conveniences. They encode deep assumptions about causality, dimensionality, and tractability, presuming that complex systems can be coordinated through linear descriptions embedded in stable contexts.

The work of Barbour, Penrose, and the RSVP framework constitutes a sustained critique of this assumption. Barbour’s denial of fundamental time replaces temporal succession with an ordering of configurations in a relational space, where change is primary and duration is a derived quantity rather than a governing parameter (Barbour 1999). Penrose’s conformal cyclic cosmology performs a complementary move with respect to scale, arguing that when rest mass becomes physically irrelevant, metric size loses meaning and only conformal relations persist at the deepest level of description (Penrose 2010). RSVP extends this line of thought by rejecting cosmic expansion itself as a primitive explanatory mechanism, treating spacetime geometry as an emergent description of irreversible entropy flow and constraint relaxation rather than as a causal arena.

Across these approaches, explanation is displaced from narrative progression to structural selection. Events are no longer explained by their position in a story but by their admissibility within a constrained configuration space. Low-entropy initial conditions, temporal asymmetry, and large-scale order are not imposed as boundary stipulations but derived from deeper invariants governing the space of possibilities. This shift marks a transition from descriptive coherence to generative adequacy, privileging frameworks that can account for why certain structures persist while others are dynamically suppressed.

The implications of this shift extend beyond physics. Narrative-based coordination functions effectively only when the relevant possibility space is narrow enough that local reasoning and symbolic compression preserve causal fidelity. As systems grow in dimensionality and interdependence, narrative explanation increasingly decouples from leverage. Inflationary cosmology, algorithmic theories of mind, and incentive-based economic storytelling persist not because they resolve underlying structural tensions, but because they preserve familiar explanatory forms. Their resilience reflects sociological path dependence rather than epistemic sufficiency (Kuhn 1962).

The replacement of narrative with relational structure does not entail relativism or explanatory retreat. On the contrary, it demands stronger commitments. A relational framework forces explicit engagement with constraints, tradeoffs, and irreversibilities. In cosmology, this appears as the demand to explain low gravitational entropy without ad hoc smoothing mechanisms (Penrose 1979). In RSVP, it appears as the insistence that redshift, structure formation, and temporal direction be derived from entropy gradients rather than postulated expansion. When this demand is applied to social and economic systems, it points inexorably toward simulation as the dominant epistemic mode, because only simulation can expose the structure of high-dimensional constraint spaces without collapsing them prematurely into narrative form.

3 Simulation as the Native Epistemology of Production

Contemporary production is increasingly organized around epistemic practices that differ fundamentally from narrative planning or market-mediated storytelling. Across advanced industrial domains, decisions are no longer made primarily through linear projections, managerial intuition, or marginal price adjustments. Instead, they are governed by simulation infrastructures that explore large configuration spaces before physical commitment occurs. Digital twins, factory simulators, supply-chain optimizers, and integrated CAD-to-CAM pipelines evaluate counterfactual trajectories under physical, logistical, and thermodynamic constraints, producing knowledge through structured exploration rather than discursive justification.

This development represents a qualitative epistemic shift. Classical economic theory presumes that coordination emerges from decentralized signals aggregated through markets, with prices functioning as compressed narratives of scarcity and preference (Smith 1776; Hayek 1945). While such mechanisms remain relevant at certain scales, they are increasingly bypassed in contexts where failure is catastrophic, coupling is tight, and path dependence dominates. Semiconductor fabrication plants, aerospace assembly lines, and large-scale energy systems are not optimized through discourse or marginal adjustment alone, but through exhaustive simulation of throughput, latency, failure modes, and resource interdependence. In these environments, narrative explanation becomes retrospective, serving to legitimize decisions already determined by constraint satisfaction.

Simulation-native production systems differ from narrative planning not merely in speed or accuracy, but in the kind of knowledge they generate. Narratives operate by suppressing dimensionality, selecting a small number of salient variables and arranging them in causal sequence. Simulations, by contrast, preserve dimensionality long enough to reveal emergent structure. They expose nonlinear interactions, threshold effects, and unintended consequences that cannot be inferred reliably from compressed descriptions. As a result, they alter the relationship between belief and commitment. One cannot meaningfully advocate a production schedule, a logistics topology, or a materials substitution without confronting its modeled consequences across time.

This asymmetry between narrative belief and simulated commitment mirrors longstanding critiques of computational reductionism in theories of mind. Penrose’s argument that genuine understanding cannot be reduced to algorithmic rule-following rests in part on the distinction between symbol manipulation and engagement with structure (Penrose 1989). Simulation-based reasoning similarly exceeds the expressive capacity of narrative description, not by invoking non-computable processes, but by refusing premature compression. It forces agents to internalize the geometry of constraint spaces rather than merely recite outcomes.

If production has already crossed this epistemic threshold, the persistence of narrative-dominated institutions elsewhere becomes increasingly untenable. Education systems that prioritize credentialed explanation over model-based reasoning, and social media platforms that optimize engagement rather than consequence visibility, increasingly resemble inflationary patches applied to failing explanatory regimes. They preserve surface coherence while deferring structural reckoning. As simulation becomes the native epistemology of production, the gap between symbolic preference expression and causal leverage widens, creating pressure for a corresponding transformation in the institutions that mediate collective choice.

4 From Discourse to Counterfactual Play

The transition to simulation-native production exposes a structural mismatch between the epistemic demands of collective decision-making and the institutions traditionally tasked with mediating it. Discourse-based mechanisms such as debate, voting, and persuasion presume that preferences are stable, communicable, and largely separable from the complex dynamics of the systems they address. These assumptions become increasingly untenable when decisions concern macrostructures whose consequences unfold over long horizons, exhibit nonlinear coupling, and generate irreversible commitments.

Preferences regarding energy transitions, technological acceleration, terraformation strategies, or ecological restoration cannot be meaningfully specified in isolation from their trajectories. Expressed symbolically, such preferences collapse high-dimensional outcome spaces into slogans or positions that obscure tradeoffs and suppress uncertainty. The result is not merely misinformation but structural misrepresentation, in which the form of expression itself distorts the domain being governed.

Narrative discourse, effective for low-dimensional coordination, becomes actively misleading when applied to systems dominated by path dependence.

A simulation-based civic substrate replaces discourse with counterfactual play. In such an environment, individuals and groups explore alternative futures within a shared, rule-governed space that encodes physical, economic, and ecological constraints. Engagement is not evaluated by assent or rhetorical alignment but by sustained interaction with trajectories and their irreversible consequences. Commitment is demonstrated through time spent navigating tradeoffs, revising strategies, and responding to emergent outcomes rather than through symbolic affirmation.

This form of participation is often dismissed as gamification, but such dismissal rests on a misunderstanding of the epistemic role of play. In scientific and engineering contexts, exploratory modeling and scenario testing are recognized as essential forms of serious work. Strategic simulation environments formalize this practice at a civic scale. They compel participants to internalize growth curves, feedback loops, and constraint surfaces in a manner analogous to the way laboratory experiments force theorists to confront recalcitrant phenomena. What matters is not the fictional veneer of the environment but the rigor of its constraints and the persistence of its state.

Historically, long-horizon strategy simulations have already demonstrated this capacity in limited domains. Their significance lies not in their subject matter but in their structural form, which forces reasoning about futures as evolving configurations rather than as endpoints. When grounded in real physical and institutional constraints, such environments become instruments for preference revelation rather than entertainment. Participants discover what they are willing to trade, protect, or abandon only by living through the consequences of their choices across extended trajectories.

This transition mirrors the movement in cosmology away from narrative expansion histories toward conformal and entropic descriptions. Just as the universe is no longer explained by a story of smooth growth alone, societies can no longer be coordinated by stories of progress, disruption, or innovation detached from structural feasibility. Simulation becomes the medium through which legitimacy is earned, because it is the medium in which consequences are made unavoidable.

5 Simulation, Preference Revelation, and Macrostructural Selection

A simulation-centered civic infrastructure does not merely improve the information available to decision-makers; it alters the very ontology of preference. In narrative-based systems, preferences are treated as antecedent facts that can be aggregated, negotiated, or represented symbolically. They are assumed to exist prior to engagement with the systems they govern. In a simulation-based system, preferences are instead revealed through interaction with structured possibility spaces. What agents claim to want is repeatedly revised as they encounter constraints, unintended consequences, and long-term dynamics that were invisible at the level of discourse.

This distinction parallels developments in economic theory concerning revealed preference, but extends them into domains where outcomes cannot be decomposed into isolated choices. Classical revealed preference theory infers preferences from observed selections among static bundles (Samuelson 1938). Simulation-based preference revelation infers commitments from sustained engagement with trajectories. What matters is not a single choice but the willingness to continue inhabiting a path as its consequences unfold. Preferences become functionals over trajectories rather than orderings over outcomes.

In this sense, simulation environments function as instruments of macrostructural selection. Competing energy regimes, technological acceleration profiles, terraformation strategies, or governance architectures can be instantiated as parameterized worlds whose evolution reflects both physical

limits and institutional design. Selection among these worlds does not occur through rhetorical victory or ideological coherence, but through demonstrated viability under shared constraints. Trajectories that collapse under their own assumptions lose adherents not because they are argued against, but because their consequences become intolerable when explored in full.

The analogy to cosmology is more than metaphorical. Penrose's insistence that the Big Bang's low entropy must be explained by inherited conformal structure rather than inflationary smoothing reflects a demand that global order be justified by deeper invariants rather than imposed boundary conditions (Penrose 2010). RSVP extends this demand by locating temporal direction and apparent expansion in entropy gradients rather than in metric growth. A simulation-based society applies the same demand to its own futures. Macrostructures must be shown to arise naturally from constraints and interactions, not from declarations of intent, moral certainty, or narrative appeal.

This reconceptualization of preference has immediate implications for education. Education conceived as narrative transmission trains individuals to repeat explanations without internalizing the dynamics that make those explanations operative. In contrast, education embedded within shared simulation environments trains individuals to reason in terms of models, sensitivities, and path dependence. Competence becomes inseparable from the ability to navigate and modify structured systems, just as scientific understanding becomes inseparable from the ability to manipulate formal representations. Expertise is no longer defined by credentialed knowledge alone, but by demonstrated capacity to sustain viable trajectories over time (Collins and Evans 2007).

The same logic renders contemporary social media structurally obsolete. Platforms optimized for attention and engagement amplify symbolic alignment while suppressing consequence visibility. They allow preferences to be expressed without cost and identities to be performed without commitment. In a simulation-centered civic infrastructure, identity and trust attach instead to maintained trajectories. Reputation emerges from demonstrated stewardship of complex systems rather than from visibility or rhetorical dominance. Preference revelation becomes inseparable from responsibility, because to express a preference is to accept the burden of its simulated consequences.

6 Irreversibility, Legitimacy, and the End of Feed-Based Coordination

The decisive epistemic advantage of simulation over narrative lies in its treatment of irreversibility. Narratives can be revised without cost, overwritten by new framings, or abandoned when they lose persuasive force. Simulation, by contrast, accumulates history. Once a trajectory is entered and explored, its consequences cannot be unseen, and its costs cannot be retroactively erased. This accumulation of consequence aligns simulation-based coordination with the thermodynamic structure emphasized in RSVP, where time is not a parameter indexing events but a record of irreversible entropy production.

Legitimacy in such a system no longer derives primarily from procedural assent or rhetorical success. It arises from demonstrated navigation of irreversible processes. Trajectories earn legitimacy insofar as they remain viable under sustained exposure to constraint, uncertainty, and perturbation. This conception of legitimacy is neither technocratic nor purely procedural. It is epistemic in the sense that it privileges demonstrated understanding of structure over symbolic conformity, and moral in the sense that it binds preference expression to consequence-bearing commitment.

This reconceptualization renders feed-based coordination structurally obsolete. Social media feeds flatten time into an endless present, erasing consequence through constant refresh. They privilege immediacy, alignment, and reaction over accumulation and stewardship. In doing so, they function as anti-thermodynamic systems, continuously resetting attention without allowing history

to bind future action. Such systems are well suited to narrative persuasion but fundamentally misaligned with the coordination of irreversible processes.

Simulation environments, by contrast, enforce temporal thickness. Decisions persist, compound, and constrain future options. Early commitments shape later possibilities, and recovery from error is costly rather than cosmetic. This persistence mirrors physical reality more closely than any narrative medium. As a result, simulation-based coordination produces a qualitatively different form of political engagement. Participants cannot simply endorse outcomes; they must inhabit paths. Disagreement does not vanish, but it becomes structured around alternative trajectories rather than rival stories.

The replacement of feeds with simulations also resolves a longstanding tension in democratic theory. Collective choice mechanisms struggle when outcomes are highly path-dependent and nonlinear, as is the case with climate intervention, large-scale automation, or planetary engineering. Voting on such issues abstracts away precisely the dynamics that determine success or failure. Simulation-based participation allows divergence rather than forcing premature convergence. Multiple futures can be explored in parallel, with resources, attention, and legitimacy gradually flowing toward those that demonstrate robustness under shared constraints.

This does not imply a technocratic substitution of optimization for politics. On the contrary, it expands the space of meaningful political action. Values enter the system not as slogans but as constraints, protected invariants, or cost functions whose implications can be examined rather than asserted. Ethical commitments become parameters that shape trajectories, making tradeoffs explicit without reducing them to utility maximization. In this sense, simulation does not eliminate disagreement; it makes disagreement legible at the level of structure rather than at the level of identity.

7 Simulation Civilizations and the Reorganization of Collective Agency

The transition from narrative coordination to simulation-based coordination marks a qualitative shift in the structure of collective agency. In narrative civilizations, agency is exercised primarily through persuasion, legitimacy is conferred through recognition, and action is mediated by institutions that translate symbolic alignment into force. Collective outcomes depend on the ability of narratives to mobilize belief and suppress complexity. In simulation civilizations, agency is exercised through the exploration, maintenance, and revision of trajectories. Legitimacy emerges from demonstrated viability under constraint, and action proceeds through the gradual scaling of configurations that have proven robust within shared models.

This shift resolves a persistent tension in both political economy and philosophy of science. Large-scale systems increasingly exhibit dynamics that defeat intuition, local reasoning, and linear causal inference. Climate systems, technological ecosystems, and automated economies display threshold effects, hysteresis, and long-range coupling that render narrative explanation unreliable. Attempts to govern such systems through discourse alone inevitably lag behind the dynamics they seek to control, producing cycles of overconfidence, crisis, and retrospective rationalization. Simulation does not eliminate uncertainty, but it localizes uncertainty within structured spaces where sensitivity can be tested rather than asserted.

The reorganization of collective agency around simulation also reframes the meaning of planning. Twentieth-century planning failures are often cited as evidence that large-scale foresight is impossible or inherently authoritarian. Yet these failures were less a consequence of ambition than of epistemic limitation. Planners lacked the computational and representational tools required to explore

high-dimensional counterfactual spaces, and thus substituted rigid schemas for genuine constraint navigation. Contemporary factory simulators demonstrate that this limitation no longer holds. The same epistemic machinery that allows firms to design complex production systems can, in principle, be generalized to social-scale questions without collapsing into centralized command.

The analogy to cosmology remains instructive. Penrose’s critique of inflation is not that it produces empirically incorrect predictions, but that it displaces explanation from deep structure to contingent dynamics (Penrose 2010). RSVP’s critique of expansion similarly targets explanatory shortcuts that obscure entropy-driven organization. In the civic domain, narrative politics plays an analogous role, offering coherence without constraint. Simulation-based governance restores explanation by forcing alignment between intent and consequence, replacing rhetorical plausibility with structural admissibility.

Under this reorganization, collective agency becomes less about deciding what to believe and more about deciding what to sustain. Futures are not chosen by vote alone but by the willingness of participants to inhabit and develop trajectories whose costs and risks are explicit. Political disagreement shifts from disputes over symbols to divergence in modeled commitments. Coordination emerges not through consensus but through selective convergence on futures that remain viable under continued exploration. In this sense, simulation civilizations do not abolish politics; they relocate it to the space where structure, consequence, and responsibility intersect.

8 Architecture of Civic Simulation Systems

A civic simulation infrastructure must be understood not as a single monolithic model, but as a layered architecture designed to accommodate heterogeneous forms of participation, uncertainty, and scale. At its core lies a set of formally specified constraint models encoding physical, ecological, economic, and institutional limits. These models must be explicit, inspectable, and revisable, since their assumptions constitute the normative and epistemic substrate of the system. Unlike narrative institutions, which obscure their assumptions behind rhetoric or tradition, simulation-based systems surface their commitments as parameters and equations.

Above this core sits a trajectory engine responsible for propagating states forward under constraint. This engine must support branching, rollback for analysis, and persistence for commitment. Crucially, rollback is an analytic operation rather than a political one. While users may explore alternative branches, trajectories that become socially instantiated accumulate irreversible history. This distinction preserves exploratory freedom without collapsing consequence.

Interface layers mediate access to this structure. A civic simulation cannot presume uniform technical expertise. It must therefore support multiple representational modalities, ranging from high-level narrative summaries to direct parameter manipulation and full model inspection. These modalities are not separate systems but projections of a single underlying state space, ensuring that accessibility does not entail epistemic fragmentation.

Verification and validation occupy a distinct architectural role. Unlike predictive simulations optimized for accuracy alone, civic simulations must be optimized for robustness under disagreement. Validation therefore concerns not only empirical fidelity but sensitivity transparency. Participants must be able to see how outcomes depend on assumptions, which parameters dominate uncertainty, and where models are fragile. This requirement distinguishes civic simulation from technocratic forecasting and aligns it with democratic legitimacy grounded in intelligibility rather than authority.

Finally, governance of the simulation itself must be endogenous. Rules for modifying constraints, introducing new modules, or deprecating outdated structures must be part of the simulated institutional space rather than external interventions. This recursive design ensures that civic

simulation remains adaptive without reverting to narrative fiat, preserving alignment between epistemic form and political function.

9 Historical Precedents and Failed Transitions

The proposal that simulation could function as civic infrastructure inevitably invites comparison with earlier attempts to replace narrative or market coordination with formalized planning and systems analysis. Twentieth-century history provides several instructive failures, most notably in large-scale state planning, military systems analysis, and early cybernetic governance projects. These failures are often cited as evidence that simulation-based coordination is inherently technocratic, brittle, or politically illegitimate. Such conclusions, however, conflate epistemic ambition with epistemic capacity.

Centralized planning efforts such as those undertaken by Gosplan relied on static models, limited computational power, and coarse abstractions that could not accommodate nonlinear dynamics or rapid feedback. Their failure was not due to the use of formal models per se, but to the inability of those models to represent high-dimensional constraint spaces or to evolve in response to unanticipated interactions. Similarly, mid-century systems analysis in military and policy contexts often collapsed complex social dynamics into simplified optimization targets, producing brittle strategies that failed outside narrowly defined conditions.

What distinguishes contemporary simulation-native systems from these historical precedents is not merely increased computational capacity, but a shift in epistemic orientation. Modern simulations are exploratory rather than prescriptive. They are designed to expose sensitivity, failure modes, and uncertainty rather than to produce singular optimal plans. This orientation aligns simulation with experimental science rather than command-and-control governance. The simulation does not dictate outcomes; it structures the space in which outcomes can be meaningfully explored.

The lesson of historical failure is therefore not that simulation is incompatible with civic coordination, but that simulation cannot substitute for politics by collapsing value conflicts into optimization targets. Civic simulation, as proposed here, does not eliminate disagreement or pluralism. Instead, it provides a medium in which disagreement can be expressed as divergence among trajectories rather than as conflict over narratives. By learning from past failures, simulation civilizations can avoid repeating the mistake of treating models as substitutes for judgment rather than as instruments for its exercise.

10 Democratic Legitimacy in Simulation-Based Systems

Any proposal to reorganize civic coordination around simulation must confront the question of democratic legitimacy. Traditional democratic theory grounds legitimacy in procedures of participation, deliberation, and consent, often emphasizing equality of voice and transparency of reasoning. Simulation-based systems appear, at first glance, to threaten these ideals by privileging technical expertise and embedding decisions within complex models that are not immediately intelligible to all participants. Addressing this concern requires a reconceptualization of legitimacy suited to high-dimensional, path-dependent systems.

Legitimacy in simulation-based systems cannot rest solely on procedural equality in the expression of opinion, because opinions detached from consequences lack epistemic weight. Instead, legitimacy emerges from the availability of meaningful participation in consequence-bearing exploration. Participants need not possess identical technical skills, but they must have access to representations that allow them to understand how trajectories respond to constraints and interventions. Equality

is preserved not by uniformity of expertise, but by symmetry of access to the underlying structure and to the capacity to contest its assumptions.

Deliberation does not disappear in such systems; it changes form. Rather than arguing primarily about desired endpoints, participants deliberate about constraints, assumptions, and tolerable risks. Public reasoning shifts from the exchange of justifications to the joint interrogation of models. This form of deliberation satisfies the core demand of democratic publicity, since the reasons for outcomes are embedded in visible structures rather than hidden in authority or rhetoric.

A common objection holds that simulation-based participation advantages technical elites and risks reproducing existing power asymmetries. This objection is serious, but it applies with equal or greater force to contemporary institutions whose technical substrates are already opaque, such as financial markets, algorithmic governance systems, and platform economies. Simulation-based civic infrastructure, by contrast, makes its technical commitments explicit and contestable. Expertise becomes accountable to structure rather than mystified through institutional prestige.

Legitimacy in this context is therefore epistemic rather than merely procedural. It depends on whether participants can trace outcomes back to assumptions, understand the sources of disagreement, and modify constraints through shared processes. Simulation does not replace democracy with optimization; it redefines democratic participation as engagement with the structures that generate futures rather than as assent to narratives about them.

11 Value Representation and Normative Constraints

A central challenge for any simulation-based civic infrastructure is the representation of values without collapsing them into hidden technical assumptions. Critics of model-driven governance have long noted that simulations inevitably encode normative commitments, whether explicitly or implicitly, through their choice of variables, constraints, and evaluation criteria. The question is therefore not whether values enter simulation, but how they are made legible, contestable, and revisable within a shared civic framework.

In narrative institutions, values are typically expressed symbolically through slogans, platforms, or moral claims whose practical implications remain underdetermined. Simulation forces a different mode of expression. Values appear as constraints on admissible trajectories, as protected invariants that cannot be violated without terminating a branch, or as costs that must be explicitly borne when tradeoffs are made. This translation does not eliminate ethical disagreement, but it relocates it from the level of rhetoric to the level of structural consequence.

This relocation has important epistemic effects. When values are encoded as constraints, their interaction with physical and social limits becomes visible. Participants can observe not only whether a value is affirmed, but what it excludes, delays, or destabilizes. Ethical commitments thus acquire a dynamic character, revealing tensions among goals that narrative discourse often suppresses. Disagreement becomes productive rather than polarizing, because it manifests as divergence among trajectories rather than as mutual negation.

Concerns about technocratic capture arise when value encoding is treated as a one-time design decision rather than an ongoing civic process. A simulation-based infrastructure must therefore include mechanisms for revising normative constraints through participation that is itself modeled and persistent. Changes to value representations accumulate history and generate consequences, preventing values from being rewritten opportunistically without acknowledgment of their effects. This persistence distinguishes civic simulation from both technocratic imposition and purely symbolic ethics.

Importantly, the explicit encoding of values does not imply moral reductionism. Not all values

need to be quantifiable, and not all constraints need to be commensurable. Some commitments may function as categorical exclusions rather than optimization targets. The key requirement is not numerical precision, but structural clarity. Values must shape the space of possible futures in ways that participants can inspect, challenge, and inhabit.

By making normative commitments explicit and consequence-bearing, simulation-based civic infrastructure restores a connection between ethics and action that narrative coordination erodes. Moral claims cease to function as identity markers or persuasive devices and instead become architectural features of shared futures. In this way, simulation does not neutralize values; it gives them durable form.

12 Objections, Limits, and Failure Modes

No proposal for simulation-based civic infrastructure is complete without addressing its potential failure modes. Simulation does not confer epistemic immunity, and its misuse can reproduce or amplify the very pathologies it is intended to correct. A rigorous account must therefore confront the strongest objections directly, not as rhetorical hurdles, but as structural risks intrinsic to the medium.

One such risk is the reality gap objection. All simulations are abstractions, and abstraction necessarily excludes detail. If the excluded details systematically advantage certain groups or assumptions, simulation can entrench bias under the guise of objectivity. This risk cannot be eliminated, but it can be mitigated by enforcing transparency about model scope, uncertainty, and sensitivity. Civic simulation must foreground not only outcomes but also where and why the model may fail. Legitimacy depends on participants understanding the limits of representation, not on believing in its completeness.

A second objection concerns manipulability. Any structured system can be gamed, and simulations are no exception. Participants may learn to exploit model artifacts rather than engage substantively with constraints. However, this risk is not unique to simulation; it is endemic to all institutional systems, including markets, elections, and bureaucracies. The distinguishing feature of simulation is that exploitability itself becomes visible. When trajectories diverge sharply under slight parameter changes, or when success depends on brittle assumptions, such fragility can be exposed and corrected through model revision rather than concealed behind narrative justification.

A third concern involves exclusion and cognitive inequality. Simulation literacy may become a new form of gatekeeping, privileging those with technical training or cognitive resources. This danger mirrors historical patterns in which new epistemic forms initially concentrate power. Addressing it requires deliberate architectural design that supports progressive engagement, multiple representational layers, and collective scaffolding. Accessibility is not achieved by simplifying models beyond usefulness, but by enabling participants to build understanding over time through interaction.

Finally, there is the risk of overreach. Simulation is not a substitute for lived experience, local knowledge, or moral judgment. Certain domains resist formalization, and attempts to simulate them exhaustively may produce false confidence. A simulation-based civic infrastructure must therefore coexist with, rather than eliminate, other forms of knowledge. Its role is to structure exploration where complexity overwhelms narrative reasoning, not to colonize all domains of human judgment.

These objections do not invalidate the core thesis. They clarify its conditions of success. Simulation becomes a civic asset only when it is treated as a fallible, contestable, and evolving medium rather than as an oracle. When its limits are acknowledged and incorporated into governance, simulation can discipline power rather than concentrate it. When its limits are denied, it becomes merely another narrative, albeit a more technical one.

13 Conclusion

The argument developed in this essay situates the rise of factory simulation, digital twins, and counterfactual optimization within a broader intellectual movement that rejects narrative primitives in favor of relational, constraint-driven structure. From Barbour's denial of fundamental time to Penrose's elimination of absolute scale at cosmological boundaries and RSVP's treatment of spacetime as an emergent record of irreversible entropy flow, the trajectory is consistent. Explanation advances by abandoning stories told against fixed backgrounds and embracing models that expose the geometry of possibility spaces and the costs of irreversibility.

Applied to social and economic systems, this trajectory implies that discourse-centric institutions have reached their epistemic limits. As production becomes simulation-native, legitimacy migrates away from institutions optimized for persuasion, engagement, and symbolic alignment toward environments in which futures are explored rather than proclaimed. Education shifts from narrative transmission to model navigation. Social media yields to shared simulation spaces where preferences are revealed through sustained engagement with consequences. Collective choice moves from rhetorical aggregation toward macrostructural selection among trajectories that demonstrate robustness under explicit constraints.

This claim is not utopian, nor does it assume that simulation will produce consensus or eliminate conflict. On the contrary, simulation foregrounds disagreement by making tradeoffs unavoidable and divergence explicit. Its significance lies not in harmonization but in intelligibility. It provides a medium in which competing visions of the future can be developed, tested, and compared without collapsing complexity into slogans or identity markers. In doing so, it binds preference expression to responsibility, restoring a connection between belief and consequence that narrative coordination systematically erodes.

The emergence of simulation as civic infrastructure is therefore not a matter of cultural preference or technological fashion. It is a structural response to the increasing dimensionality, path dependence, and irreversibility of the systems societies must now govern. Just as cosmology could not remain narrative once confronted with entropy, singularities, and scale invariance, societies cannot remain narrative once confronted with planetary engineering, automated production, and long-horizon technological trajectories. The replacement of narrative coordination with simulation-based preference revelation marks a necessary reconfiguration of collective agency, one in which futures are not merely imagined, but inhabited, tested, and selectively sustained.

Appendices

A Trajectory Spaces and Preference Functionals

Let \mathcal{S} be a state space equipped with a σ -algebra Σ and a constraint operator

$$\mathcal{C} : \mathcal{S} \rightarrow \{0, 1\},$$

where $\mathcal{C}(s) = 1$ indicates admissibility.

Define a trajectory as a map

$$\gamma : [0, T] \rightarrow \mathcal{S}$$

such that $\mathcal{C}(\gamma(t)) = 1$ for all $t \in [0, T]$.

The space of admissible trajectories is

$$\Gamma_T = \{\gamma \mid \gamma \in C([0, T], \mathcal{S}), \mathcal{C}(\gamma(t)) = 1 \forall t\}.$$

Let $\Gamma = \bigcup_{T>0} \Gamma_T$ denote the space of all finite trajectories.

A preference functional is defined as a map

$$\Pi : \Gamma \rightarrow \mathbb{R} \cup \{-\infty\},$$

where $\Pi(\gamma) = -\infty$ for any trajectory violating a protected constraint.

Two trajectories γ_1, γ_2 are said to be preference-equivalent if

$$\Pi(\gamma_1) = \Pi(\gamma_2).$$

A revealed preference ordering \succ over trajectories is induced by

$$\gamma_1 \succ \gamma_2 \iff \Pi(\gamma_1) > \Pi(\gamma_2).$$

Preference convergence is defined as stabilization of $\Pi(\gamma_t)$ under trajectory extension:

$$\lim_{T \rightarrow \infty} \Pi(\gamma|_{[0, T]}) \text{ exists.}$$

B Constraint Manifolds and Admissibility Geometry

Let \mathcal{S} be a smooth manifold representing the instantaneous configuration space of a system. Let

$$\Phi : \mathcal{S} \rightarrow \mathbb{R}^k$$

be a constraint map encoding k independent constraints.

Define the admissible manifold

$$\mathcal{M} = \{s \in \mathcal{S} \mid \Phi(s) = 0\}.$$

Assume $\text{rank}(D\Phi) = k$ on \mathcal{M} so that \mathcal{M} is a smooth embedded submanifold of codimension k . Let $T_s\mathcal{M}$ denote the tangent space at $s \in \mathcal{M}$. Admissible velocities $v \in T_s\mathcal{S}$ satisfy

$$D\Phi(s) \cdot v = 0.$$

Define a constraint metric

$$g_C = g_S + \lambda D\Phi^\top D\Phi$$

for $\lambda > 0$, where $g_{\mathcal{S}}$ is a background metric on \mathcal{S} .

Admissible trajectories are geodesics of g_C restricted to \mathcal{M} .

Define the constraint curvature tensor \mathcal{K} via the second fundamental form

$$\mathcal{K}(u, v) = (\nabla_u v)^\perp$$

for $u, v \in T\mathcal{M}$.

Regions of high $\|\mathcal{K}\|$ correspond to sensitive dependence of admissibility under perturbation.

A constraint bifurcation occurs at $s^* \in \mathcal{M}$ when

$$\det(D^2\Phi(s^*)) = 0.$$

Such points separate topologically distinct admissible regions in \mathcal{M} .

C Irreversibility, Entropy Accumulation, and Trajectory Ordering

Let Γ denote the space of admissible trajectories defined over intervals $[0, T]$ with values in an admissible manifold \mathcal{M} .

Define an entropy functional

$$S : \Gamma \rightarrow \mathbb{R}$$

such that for any trajectory γ and any $0 \leq t_1 < t_2 \leq T$,

$$S(\gamma|_{[0, t_2]}) \geq S(\gamma|_{[0, t_1]}).$$

A trajectory γ is said to be reversible on $[t_1, t_2]$ if

$$S(\gamma|_{[0, t_2]}) = S(\gamma|_{[0, t_1]}),$$

and irreversible otherwise.

Define the entropy production rate

$$\sigma(t) = \frac{d}{dt} S(\gamma|_{[0, t]}),$$

assumed to satisfy $\sigma(t) \geq 0$ almost everywhere.

Introduce a preorder \preceq on Γ by

$$\gamma_1 \preceq \gamma_2 \iff S(\gamma_1) \leq S(\gamma_2)$$

for trajectories defined over equal durations.

The preorder induces equivalence classes

$$[\gamma] = \{\gamma' \mid S(\gamma') = S(\gamma)\}.$$

Define an irreversible extension operator

$$\mathcal{E}_\Delta : \Gamma_T \rightarrow \Gamma_{T+\Delta}$$

such that

$$S(\mathcal{E}_\Delta(\gamma)) > S(\gamma)$$

for all $\Delta > 0$.

A trajectory is said to be terminally irreversible if no admissible inverse extension exists:

$$\nexists \mathcal{E}_{-\Delta} \text{ such that } \mathcal{E}_{-\Delta}(\gamma) \in \Gamma_{T-\Delta}.$$

Define a partial order on equivalence classes by

$$[\gamma_1] \prec [\gamma_2] \iff S(\gamma_1) < S(\gamma_2),$$

which induces a directed acyclic structure on the quotient space Γ / \sim .

This ordering defines a temporal orientation intrinsic to trajectory space, independent of any external time parameter.

D Branching Dynamics, Commitment Operators, and Preference Collapse

Let Γ_T denote the space of admissible trajectories defined on $[0, T]$.

Define a branching operator

$$\mathcal{B} : \Gamma_T \rightarrow \mathcal{P}(\Gamma_{T+\Delta})$$

such that

$$\mathcal{B}(\gamma) = \{\gamma' \mid \gamma'|_{[0, T]} = \gamma\},$$

where $\Delta > 0$ is fixed.

Let $\mathcal{B}(\gamma)$ be finite or countable.

Define a commitment operator

$$\mathcal{K} : \mathcal{P}(\Gamma_{T+\Delta}) \rightarrow \Gamma_{T+\Delta}$$

such that

$$\mathcal{K}(\mathcal{B}(\gamma)) = \gamma^*,$$

where γ^* is selected and all other branches are rendered inaccessible.

Define commitment irreversibility by

$$\nexists \gamma' \in \mathcal{B}(\gamma) \setminus \{\gamma^*\}$$

after application of \mathcal{K} .

Let $\Pi : \Gamma \rightarrow \mathbb{R} \cup \{-\infty\}$ be a preference functional.

Define preference dispersion at time T by

$$D_T = \sup_{\gamma_1, \gamma_2 \in \mathcal{B}(\gamma|_{[0, T-\Delta]})} |\Pi(\gamma_1) - \Pi(\gamma_2)|.$$

Preference collapse occurs at T if

$$\lim_{T \rightarrow \infty} D_T = 0.$$

A system is said to exhibit structural preference revelation if

$$\exists T^* \text{ such that } \forall T > T^*, D_T < \epsilon$$

for arbitrarily small $\epsilon > 0$.

Under repeated application of \mathcal{B} and \mathcal{K} , the admissible trajectory space contracts onto a subset

$$\Gamma^* \subset \Gamma$$

such that all remaining trajectories are preference-equivalent.

This contraction defines a fixed-point structure in trajectory space corresponding to stabilized macrostructural commitment.

E Macrostructural Fixed Points and Stability Classes

Let Γ denote the space of admissible trajectories modulo preference equivalence, endowed with the partial order induced by entropy accumulation and commitment operators.

Define a macrostructure as an equivalence class

$$\mathcal{A} \subset \Gamma$$

such that all $\gamma \in \mathcal{A}$ are preference-equivalent and terminally irreversible.

Introduce a transition operator

$$\mathcal{T} : \Gamma \rightarrow \Gamma$$

representing one full cycle of branching, evaluation, and commitment.

A macrostructure \mathcal{A} is a fixed point of \mathcal{T} if

$$\mathcal{T}(\gamma) \in \mathcal{A} \quad \forall \gamma \in \mathcal{A}.$$

Define the basin of attraction of \mathcal{A} as

$$\mathcal{B}(\mathcal{A}) = \{\gamma \in \Gamma \mid \lim_{n \rightarrow \infty} \mathcal{T}^n(\gamma) \in \mathcal{A}\}.$$

Two macrostructures \mathcal{A}_1 and \mathcal{A}_2 are said to be dynamically distinguishable if

$$\mathcal{B}(\mathcal{A}_1) \cap \mathcal{B}(\mathcal{A}_2) = \emptyset.$$

Define a stability functional

$$\Lambda(\mathcal{A}) = \inf_{\gamma \in \mathcal{A}} \sup_{\delta\gamma} [S(\mathcal{T}(\gamma + \delta\gamma)) - S(\mathcal{T}(\gamma))],$$

where $\delta\gamma$ ranges over admissible perturbations.

A macrostructure is stable if

$$\Lambda(\mathcal{A}) \leq 0,$$

and unstable otherwise.

Let \mathfrak{M} denote the set of all macrostructural fixed points. The partially ordered set

$$(\mathfrak{M}, \prec)$$

induced by entropy ordering defines a hierarchy of viable futures.

Selection among macrostructures corresponds to convergence of collective trajectories into a particular basin of attraction under repeated application of \mathcal{T} .

F Collective Coupling, Coherence, and Multi-Agent Trajectories

Let $\{\Gamma^{(i)}\}_{i=1}^N$ denote the trajectory spaces associated with N agents, each defined over a common admissible manifold \mathcal{M} .

Define the joint trajectory space as the Cartesian product

$$\Gamma^{(N)} = \prod_{i=1}^N \Gamma^{(i)}.$$

Let

$$\gamma^{(N)}(t) = (\gamma^{(1)}(t), \dots, \gamma^{(N)}(t))$$

denote a collective trajectory.

Introduce an interaction operator

$$\mathcal{I} : \Gamma^{(N)} \rightarrow \mathbb{R}_{\geq 0}$$

measuring coupling strength among agents along a trajectory.

Assume \mathcal{I} satisfies symmetry and non-negativity:

$$\mathcal{I}(\gamma^{(N)}) = \mathcal{I}(\pi(\gamma^{(N)})), \quad \mathcal{I} \geq 0$$

for any permutation π of agent indices.

Define collective entropy as

$$S^{(N)}(\gamma^{(N)}) = \sum_{i=1}^N S(\gamma^{(i)}) + \alpha \mathcal{I}(\gamma^{(N)})$$

for coupling coefficient $\alpha \geq 0$.

A collective trajectory is admissible if

$$S^{(N)}(\gamma^{(N)}|_{[0,t]}) \text{ is non-decreasing in } t.$$

Define coherence at time T by

$$C_T = \inf_{\gamma^{(N)} \in \Gamma_T^{(N)}} \left[\sup_{i,j} d(\gamma^{(i)}(T), \gamma^{(j)}(T)) \right],$$

where d is a metric on \mathcal{M} .

A collective is said to be coherent if

$$\lim_{T \rightarrow \infty} C_T < \infty.$$

Define a collective preference functional

$$\Pi^{(N)}(\gamma^{(N)}) = \sum_{i=1}^N \Pi(\gamma^{(i)}) - \beta \mathcal{I}(\gamma^{(N)})$$

for $\beta \geq 0$.

Collective preference collapse occurs if

$$\lim_{T \rightarrow \infty} \sup_{\gamma_1^{(N)}, \gamma_2^{(N)} \in \mathcal{B}(\gamma^{(N)}|_{[0, T-\Delta])}} \left| \Pi^{(N)}(\gamma_1^{(N)}) - \Pi^{(N)}(\gamma_2^{(N)}) \right| = 0.$$

A collective macrostructure is defined as a fixed point of the induced transition operator on $\Gamma^{(N)}$ under coupling, entropy accumulation, and commitment.

Such macrostructures characterize coherent, stable collective futures emergent from interacting agent trajectories.

G Uncertainty Propagation and Robustness Criteria

Let Γ denote the space of admissible trajectories on the manifold \mathcal{M} , and let Θ be a parameter space encoding uncertain model inputs.

Assume a parametrized family of dynamics

$$\mathcal{D}_\theta : \Gamma \rightarrow \Gamma$$

for $\theta \in \Theta$.

Define a probability measure μ on Θ representing epistemic uncertainty.

For a trajectory $\gamma \in \Gamma$, define the uncertainty envelope

$$\mathcal{U}(\gamma) = \{\mathcal{D}_\theta(\gamma) \mid \theta \in \Theta\}.$$

Let $\Pi : \Gamma \rightarrow \mathbb{R} \cup \{-\infty\}$ be a preference functional. Define the worst-case valuation

$$\Pi_{\min}(\gamma) = \inf_{\theta \in \Theta} \Pi(\mathcal{D}_\theta(\gamma)),$$

and the expected valuation

$$\mathbb{E}_\mu[\Pi(\gamma)] = \int_{\Theta} \Pi(\mathcal{D}_\theta(\gamma)) d\mu(\theta).$$

Define a robustness functional

$$\mathcal{R}(\gamma) = \lambda \Pi_{\min}(\gamma) + (1 - \lambda) \mathbb{E}_\mu[\Pi(\gamma)], \quad 0 \leq \lambda \leq 1.$$

A trajectory γ^* is said to be robustly admissible if

$$\mathcal{R}(\gamma^*) \geq \mathcal{R}(\gamma) \quad \forall \gamma \in \Gamma.$$

Define robustness dominance between trajectories by

$$\gamma_1 \succ_R \gamma_2 \iff \mathcal{R}(\gamma_1) > \mathcal{R}(\gamma_2).$$

A macrostructure $\mathcal{A} \subset \Gamma$ is robust if

$$\inf_{\gamma \in \mathcal{A}} \mathcal{R}(\gamma) \geq \sup_{\gamma \notin \mathcal{A}} \mathcal{R}(\gamma).$$

Let $\mathfrak{M}_R \subset \mathfrak{M}$ denote the set of robust macrostructural fixed points.

Selection under uncertainty corresponds to convergence of collective trajectories into \mathfrak{M}_R under repeated application of the transition operator with stochastic parameter realization.

This criterion ensures that selected futures are stable not only under known constraints but under admissible uncertainty in model structure and parameters.

H Simulation Depth, Resolution Limits, and Coarse-Graining

Let \mathcal{M} be an admissible manifold equipped with a family of projections

$$\{\pi_\ell : \mathcal{M} \rightarrow \mathcal{M}_\ell\}_{\ell \in \mathbb{N}}$$

indexed by simulation resolution level ℓ , where higher ℓ corresponds to finer structural detail.

Assume each \mathcal{M}_ℓ is a quotient space of \mathcal{M} under an equivalence relation \sim_ℓ satisfying

$$s_1 \sim_\ell s_2 \implies s_1 \sim_{\ell'} s_2 \quad \forall \ell' < \ell.$$

Define the coarse-grained trajectory at level ℓ as

$$\gamma_\ell = \pi_\ell \circ \gamma.$$

Let Γ_ℓ denote the space of admissible trajectories in \mathcal{M}_ℓ .

Define a resolution-consistent preference functional

$$\Pi_\ell : \Gamma_\ell \rightarrow \mathbb{R} \cup \{-\infty\}$$

such that for any $\gamma \in \Gamma$,

$$\Pi_\ell(\gamma_\ell) \rightarrow \Pi(\gamma) \quad \text{as } \ell \rightarrow \infty.$$

Define resolution error

$$\varepsilon_\ell(\gamma) = |\Pi(\gamma) - \Pi_\ell(\gamma_\ell)|.$$

A trajectory γ is said to be resolution-stable if

$$\lim_{\ell \rightarrow \infty} \varepsilon_\ell(\gamma) = 0.$$

Define the simulation depth of a trajectory as

$$D(\gamma) = \inf\{\ell \mid \varepsilon_\ell(\gamma) < \delta\}$$

for fixed tolerance $\delta > 0$.

Let Γ^δ denote the set of trajectories admitting finite simulation depth.

Macrostructures $\mathcal{A} \subset \Gamma$ are resolution-invariant if

$$\forall \gamma \in \mathcal{A}, \quad \gamma \in \Gamma^\delta \text{ and } \sup_{\ell \geq D(\gamma)} \varepsilon_\ell(\gamma) < \delta.$$

Such macrostructures persist under coarse-graining and define futures whose preference ordering is stable across simulation resolutions.

This establishes a formal criterion for when civic-scale simulations remain meaningful despite finite computational resources and unavoidable abstraction.

I Interoperability, Modular Composition, and Constraint Gluing

Let $\{\mathcal{M}_i\}_{i \in I}$ be a family of admissible manifolds, each equipped with its own constraint map

$$\Phi_i : \mathcal{M}_i \rightarrow \mathbb{R}^{k_i}.$$

Define overlap regions $\mathcal{M}_{ij} \subset \mathcal{M}_i \times \mathcal{M}_j$ together with compatibility maps

$$\psi_{ij} : \mathcal{M}_{ij} \rightarrow \mathcal{M}_j, \quad \psi_{ji} : \mathcal{M}_{ij} \rightarrow \mathcal{M}_i,$$

satisfying cocycle conditions

$$\psi_{ik} = \psi_{jk} \circ \psi_{ij} \quad \text{on triple overlaps.}$$

The global admissible space is defined as the glued manifold

$$\mathcal{M} = \bigsqcup_{i \in I} \mathcal{M}_i / \sim$$

where $x \sim \psi_{ij}(x)$ for all $x \in \mathcal{M}_{ij}$.

Let Γ_i denote the admissible trajectory space on \mathcal{M}_i . Define the global trajectory space

$$\Gamma = \left\{ \gamma \mid \gamma|_{U_i} \in \Gamma_i \text{ and } \gamma|_{U_{ij}} \text{ is compatible under } \psi_{ij} \right\}.$$

Define local preference functionals

$$\Pi_i : \Gamma_i \rightarrow \mathbb{R} \cup \{-\infty\}.$$

A global preference functional exists if there is a map

$$\Pi : \Gamma \rightarrow \mathbb{R} \cup \{-\infty\}$$

such that

$$\Pi(\gamma) = \sum_{i \in I} \Pi_i(\gamma|_{U_i})$$

and compatibility constraints enforce

$$\Pi_i(\gamma_i) = \Pi_j(\gamma_j) \quad \text{on overlaps.}$$

Define a modular simulation as a tuple

$$(\{\mathcal{M}_i\}, \{\Phi_i\}, \{\Pi_i\}, \{\psi_{ij}\})$$

satisfying global admissibility and preference coherence.

A modular macrostructure is a fixed point of the induced transition operator on Γ invariant under restriction to all submodules.

This formalism defines interoperability as constraint-preserving gluing rather than interface-level compatibility, ensuring that composite simulations preserve admissibility, irreversibility, and preference ordering across scales.

J Limit Objects, Convergence, and Asymptotic Futures

Let $\{\Gamma_n\}_{n \in \mathbb{N}}$ be a directed system of admissible trajectory spaces indexed by simulation horizon, resolution, or institutional scope, with connecting maps

$$\rho_{n,m} : \Gamma_m \rightarrow \Gamma_n \quad \text{for } m \geq n,$$

satisfying $\rho_{n,n} = \text{id}$ and $\rho_{n,k} = \rho_{n,m} \circ \rho_{m,k}$ for all $k \geq m \geq n$.

Define the inverse limit trajectory space

$$\Gamma_\infty = \varprojlim \Gamma_n = \{(\gamma_n)_{n \in \mathbb{N}} \mid \rho_{n,m}(\gamma_m) = \gamma_n \ \forall m \geq n\}.$$

Let $\Pi_n : \Gamma_n \rightarrow \mathbb{R} \cup \{-\infty\}$ be compatible preference functionals such that

$$\Pi_n(\rho_{n,m}(\gamma_m)) = \Pi_m(\gamma_m) \quad \forall m \geq n.$$

Define the asymptotic preference functional

$$\Pi_\infty : \Gamma_\infty \rightarrow \mathbb{R} \cup \{-\infty\}, \quad \Pi_\infty((\gamma_n)) = \lim_{n \rightarrow \infty} \Pi_n(\gamma_n),$$

whenever the limit exists.

A sequence (γ_n) is said to be convergent if $\Pi_\infty((\gamma_n))$ exists and is finite.

Define an asymptotic future as an equivalence class

$$\mathcal{F} \subset \Gamma_\infty$$

such that all $(\gamma_n) \in \mathcal{F}$ are preference-equivalent under Π_∞ .

Introduce an asymptotic transition operator

$$\mathcal{T}_\infty : \Gamma_\infty \rightarrow \Gamma_\infty$$

induced by the compatible family $\{\mathcal{T}_n\}$ on each Γ_n .

An asymptotic future \mathcal{F} is stable if

$$\mathcal{T}_\infty((\gamma_n)) \in \mathcal{F} \quad \forall (\gamma_n) \in \mathcal{F}.$$

Let \mathfrak{F} denote the set of all stable asymptotic futures. The partially ordered set (\mathfrak{F}, \prec) induced by Π_∞ defines the terminal hierarchy of futures under unbounded simulation depth and horizon.

Convergence of collective simulation corresponds to selection of an element of \mathfrak{F} under refinement of resolution, extension of horizon, and accumulation of irreversible commitments.

K Computational Tractability, Approximation, and Feasible Selection

Let Γ denote the admissible trajectory space endowed with a metric d_Γ and preference functional Π .

Define the exact selection problem as

$$\gamma^* \in \arg \max_{\gamma \in \Gamma} \Pi(\gamma).$$

Assume that exact maximization over Γ is intractable. Introduce a family of approximation operators

$$\mathcal{A}_\epsilon : \Gamma \rightarrow \Gamma_\epsilon \subset \Gamma$$

indexed by tolerance $\epsilon > 0$, such that

$$\sup_{\gamma \in \Gamma} \inf_{\tilde{\gamma} \in \Gamma_\epsilon} d_\Gamma(\gamma, \tilde{\gamma}) \leq \epsilon.$$

Define the approximate preference functional

$$\Pi_\epsilon(\tilde{\gamma}) = \Pi(\tilde{\gamma}), \quad \tilde{\gamma} \in \Gamma_\epsilon.$$

An ϵ -optimal trajectory $\tilde{\gamma}^*$ satisfies

$$\Pi_\epsilon(\tilde{\gamma}^*) \geq \sup_{\tilde{\gamma} \in \Gamma_\epsilon} \Pi_\epsilon(\tilde{\gamma}) - \epsilon.$$

Define feasible selection consistency by

$$\lim_{\epsilon \rightarrow 0} \Pi(\tilde{\gamma}_\epsilon^*) = \sup_{\gamma \in \Gamma} \Pi(\gamma),$$

where $\tilde{\gamma}_\epsilon^*$ is any ϵ -optimal trajectory.

Introduce a complexity functional

$$\kappa(\epsilon) = \inf\{|\Gamma_\epsilon| \mid \mathcal{A}_\epsilon \text{ satisfies the approximation bound}\}.$$

A simulation architecture is tractable if

$$\kappa(\epsilon) = O(\epsilon^{-p})$$

for some finite $p > 0$.

Define a macrostructure \mathcal{A} to be computationally realizable if

$$\exists \epsilon_n \rightarrow 0 \quad \text{such that} \quad \tilde{\gamma}_{\epsilon_n}^* \in \mathcal{A} \text{ for all sufficiently large } n.$$

Thus selection among macrostructures is feasible when preference-optimal trajectories are stable under controlled approximation and bounded computational growth.

This criterion separates simulable futures from those that exist only as formal optima in intractable trajectory spaces.

L Information Flow, Observability, and Identifiability

Let (Γ, \mathcal{F}) be a measurable space of admissible trajectories, where \mathcal{F} is a σ -algebra generated by trajectory segments.

For each time $T > 0$, define the prefix projection

$$\pi_T : \Gamma \rightarrow \Gamma_T,$$

and let $\mathcal{F}_T = \pi_T^{-1}(\mathcal{B}(\Gamma_T))$ denote the induced σ -algebra of observable events up to T .

An observation operator is a measurable map

$$\mathcal{O}_T : \Gamma \rightarrow \mathcal{Y}_T,$$

where $(\mathcal{Y}_T, \mathcal{G}_T)$ is an observation space with σ -algebra \mathcal{G}_T .

Define the observable σ -algebra

$$\mathcal{H}_T = \mathcal{O}_T^{-1}(\mathcal{G}_T) \subseteq \mathcal{F}_T.$$

Let $\Pi : \Gamma \rightarrow \mathbb{R} \cup \{-\infty\}$ be a preference functional.

Two trajectories γ_1, γ_2 are observationally equivalent at horizon T if

$$\mathcal{O}_T(\gamma_1) = \mathcal{O}_T(\gamma_2).$$

Define the identifiability condition:

$$\gamma_1 \neq \gamma_2 \quad \Rightarrow \quad \exists T > 0 \text{ such that } \mathbb{E}[\Pi(\gamma_1) \mid \mathcal{H}_T] \neq \mathbb{E}[\Pi(\gamma_2) \mid \mathcal{H}_T].$$

A simulation system is preference-identifiable if the above condition holds almost surely with respect to any admissible uncertainty measure on Γ .

Define the information gain functional

$$I_T(\gamma) = \text{Var}(\Pi(\gamma)) - \text{Var}(\Pi(\gamma) \mid \mathcal{H}_T).$$

Preference revelation is asymptotically observable if

$$\lim_{T \rightarrow \infty} I_T(\gamma) > 0$$

for all non-equivalent $\gamma \in \Gamma$.

This establishes observability as a necessary condition for meaningful simulation-based preference selection.

M Strategic Misrepresentation and Incentive Compatibility

Let $\Gamma^{(i)}$ denote the admissible trajectory space associated with agent i , and let $\Pi^{(i)} : \Gamma^{(i)} \rightarrow \mathbb{R} \cup \{-\infty\}$ be the true preference functional of agent i .

Define a reported preference functional

$$\hat{\Pi}^{(i)} = R^{(i)}(\Pi^{(i)}),$$

where $R^{(i)}$ is a reporting operator chosen by agent i .

Let $\hat{\Gamma}$ denote the trajectory space induced by the aggregation of reported preferences $\{\hat{\Pi}^{(i)}\}$ through a collective selection operator \mathcal{K} .

The induced realized trajectory is

$$\gamma^* = \mathcal{K}(\{\hat{\Pi}^{(i)}\}).$$

Define the realized utility of agent i as

$$U^{(i)} = \Pi^{(i)}(\gamma^*).$$

A deviation $R^{(i)} \rightarrow \tilde{R}^{(i)}$ is profitable if

$$\Pi^{(i)}(\mathcal{K}(\tilde{R}^{(i)}(\Pi^{(i)}), \hat{\Pi}^{(-i)})) > \Pi^{(i)}(\mathcal{K}(\hat{\Pi}^{(i)}, \hat{\Pi}^{(-i)})).$$

The simulation mechanism is incentive compatible if no profitable deviation exists for any agent i .

Define trajectory-dominant strategy truthfulness by

$$R^{(i)} = \text{id} \quad \text{is dominant for all } i.$$

A weaker equilibrium notion is trajectory Nash equilibrium, where

$$\forall i, \quad U^{(i)}(R^{(i)}, R^{(-i)}) \geq U^{(i)}(\tilde{R}^{(i)}, R^{(-i)}).$$

Incentive compatibility under irreversibility is strengthened by requiring that misrepresentation induces entropy increase:

$$S(\gamma_{\text{misrep}}^*) > S(\gamma_{\text{truthful}}^*).$$

Such mechanisms penalize strategic distortion through irreversible trajectory degradation rather than explicit sanctions.

N Learning Dynamics and Belief Updating

Let Θ be a parameter space indexing uncertain model components, and let μ_0 be an initial belief measure over Θ .

For each $\theta \in \Theta$, let \mathcal{D}_θ denote the induced trajectory dynamics on Γ .

Define the belief state at iteration n as a probability measure μ_n on Θ .

Given an observed trajectory segment $\gamma|_{[0, T_n]}$, define the likelihood

$$\mathcal{L}(\theta \mid \gamma|_{[0, T_n]})$$

measurable with respect to the observation σ -algebra \mathcal{H}_{T_n} .

Belief updating proceeds via the operator

$$\mu_{n+1}(A) = \frac{\int_A \mathcal{L}(\theta \mid \gamma|_{[0, T_n]}) d\mu_n(\theta)}{\int_\Theta \mathcal{L}(\theta \mid \gamma|_{[0, T_n]}) d\mu_n(\theta)}$$

for measurable $A \subseteq \Theta$.

Define belief entropy

$$H(\mu_n) = - \int_\Theta \log \mu_n(\theta) d\mu_n(\theta).$$

Learning convergence occurs if

$$\lim_{n \rightarrow \infty} H(\mu_n) = H^*$$

for some finite H^* .

Define belief trajectory coherence by

$$\lim_{n \rightarrow \infty} \sup_{\theta \in \text{supp}(\mu_n)} d_\Gamma(\mathcal{D}_\theta(\gamma), \gamma) = 0.$$

A learning process is stable if belief convergence implies robustness convergence:

$$\mu_n \rightarrow \mu^* \quad \Rightarrow \quad \gamma_n^* \rightarrow \gamma^*,$$

where γ_n^* is the selected trajectory under μ_n .

This formalizes learning as convergence of belief measures coupled to stabilization of preferred trajectories under repeated simulation exposure.

O Symmetry, Equivalence, and Quotienting of Futures

Let Γ be the admissible trajectory space and let G be a group acting on Γ via

$$\alpha : G \times \Gamma \rightarrow \Gamma, \quad (g, \gamma) \mapsto g \cdot \gamma,$$

such that admissibility and entropy accumulation are preserved:

$$\gamma \in \Gamma \Rightarrow g \cdot \gamma \in \Gamma, \quad S(g \cdot \gamma) = S(\gamma).$$

Define trajectory equivalence under symmetry by

$$\gamma_1 \sim_G \gamma_2 \iff \exists g \in G \text{ such that } \gamma_2 = g \cdot \gamma_1.$$

Let the quotient space of futures be

$$\Gamma/G = \Gamma / \sim_G .$$

Assume the preference functional Π is G -invariant:

$$\Pi(g \cdot \gamma) = \Pi(\gamma) \quad \forall g \in G.$$

Then Π induces a well-defined functional

$$\bar{\Pi} : \Gamma/G \rightarrow \mathbb{R} \cup \{-\infty\}.$$

Define macrostructures on the quotient as equivalence classes

$$\bar{\mathcal{A}} \subset \Gamma/G$$

that are fixed points of the induced transition operator $\bar{\mathcal{T}}$.

Two macrostructures $\mathcal{A}_1, \mathcal{A}_2 \subset \Gamma$ are gauge-equivalent if

$$\mathcal{A}_2 = g \cdot \mathcal{A}_1 \quad \text{for some } g \in G.$$

Let \mathfrak{M}/G denote the set of inequivalent macrostructures modulo symmetry.

Selection among futures is said to be symmetry-reduced if convergence occurs in \mathfrak{M}/G rather than in \mathfrak{M} .

This formalism eliminates redundancy arising from relabeling, scaling, or institutional isomorphism, ensuring that distinct selected futures correspond to genuinely distinct macrostructural regimes.

P Catastrophic Boundaries, Absorbing States, and No-Return Sets

Let Γ be the admissible trajectory space with entropy functional S and preference functional Π .

Define a catastrophic boundary as a subset

$$\mathcal{X} \subset \Gamma$$

such that for any trajectory $\gamma \in \mathcal{X}$,

$$\forall \Delta > 0, \quad \Pi(\mathcal{E}_\Delta(\gamma)) = -\infty,$$

where \mathcal{E}_Δ denotes the irreversible extension operator.

A trajectory γ is catastrophic if

$$\exists T \text{ such that } \gamma|_{[0,T]} \in \mathcal{X}.$$

Define the absorbing set

$$\mathcal{A}_{\text{abs}} = \{\gamma \in \Gamma \mid \forall \Delta > 0, \mathcal{E}_\Delta(\gamma) = \gamma\},$$

where entropy accumulation satisfies

$$S(\mathcal{E}_\Delta(\gamma)) = S(\gamma).$$

Define a no-return set

$$\mathcal{N} \subset \Gamma$$

such that

$$\gamma \in \mathcal{N} \Rightarrow \nexists \gamma' \prec \gamma$$

under the entropy-induced partial order.

Define the catastrophe operator

$$\mathcal{C} : \Gamma \rightarrow \{0, 1\}$$

where $\mathcal{C}(\gamma) = 1$ if γ intersects \mathcal{X} .

A preference functional is catastrophe-aware if

$$\Pi(\gamma) = -\infty \quad \forall \gamma \text{ such that } \mathcal{C}(\gamma) = 1.$$

Define the safe admissible space

$$\Gamma_{\text{safe}} = \Gamma \setminus \mathcal{X}.$$

A macrostructure \mathcal{A} is viable if

$$\mathcal{A} \subset \Gamma_{\text{safe}} \quad \text{and} \quad \mathcal{T}(\mathcal{A}) \subset \Gamma_{\text{safe}}.$$

Selection dynamics constrained to Γ_{safe} ensure that catastrophic trajectories are excluded prior to preference comparison, enforcing existential constraints independently of optimization or aggregation.

This formalizes irreversible existential risk as a boundary condition in trajectory space rather than as a tradeable cost.

Q Continuum Limits, Discretization Independence, and Renormalization

Let $\{\Gamma_{\Delta t}\}_{\Delta t > 0}$ be a family of admissible trajectory spaces indexed by time discretization scale Δt , where each $\Gamma_{\Delta t}$ consists of piecewise-constant or piecewise-smooth trajectories with timestep Δt .

Assume refinement maps

$$r_{\Delta t, \Delta t'} : \Gamma_{\Delta t} \rightarrow \Gamma_{\Delta t'} \quad \text{for } \Delta t' < \Delta t$$

satisfying consistency conditions

$$r_{\Delta t', \Delta t''} \circ r_{\Delta t, \Delta t'} = r_{\Delta t, \Delta t''}.$$

Define the continuum trajectory space as the direct limit

$$\Gamma_0 = \varinjlim_{\Delta t \rightarrow 0} \Gamma_{\Delta t}.$$

Let $\Pi_{\Delta t} : \Gamma_{\Delta t} \rightarrow \mathbb{R} \cup \{-\infty\}$ be discretized preference functionals.

Discretization consistency requires

$$\lim_{\Delta t \rightarrow 0} \Pi_{\Delta t}(r_{\Delta t, 0}(\gamma_{\Delta t})) = \Pi_0(\gamma) \quad \forall \gamma \in \Gamma_0.$$

Define discretization error

$$\delta_{\Delta t}(\gamma) = |\Pi_0(\gamma) - \Pi_{\Delta t}(r_{0, \Delta t}(\gamma))|.$$

A preference ordering is discretization-independent if

$$\lim_{\Delta t \rightarrow 0} \sup_{\gamma_1, \gamma_2} |(\Pi_{\Delta t}(\gamma_1) - \Pi_{\Delta t}(\gamma_2)) - (\Pi_0(\gamma_1) - \Pi_0(\gamma_2))| = 0.$$

Introduce a renormalization operator

$$\mathcal{R}_\lambda : \Gamma_0 \rightarrow \Gamma_0$$

representing coarse-graining by factor $\lambda > 1$.

A macrostructure \mathcal{A} is renormalization-invariant if

$$\mathcal{R}_\lambda(\mathcal{A}) = \mathcal{A} \quad \forall \lambda > 1.$$

Let $\mathfrak{M}_{\text{RG}} \subset \mathfrak{M}$ denote the set of renormalization-invariant macrostructures.

These macrostructures define futures whose preference ordering, stability, and admissibility persist under refinement of temporal resolution, agent granularity, and simulation timestep, ensuring that selection is not an artifact of discretization.

This establishes continuum robustness as a necessary condition for civic-scale simulation validity.

R Domain Instantiation via Parametric Constraint Specialization

Let \mathcal{M} be a global admissible manifold equipped with constraint map

$$\Phi : \mathcal{M} \rightarrow \mathbb{R}^k.$$

Let \mathcal{D} be a domain index set. For each domain $d \in \mathcal{D}$, define a specialization operator

$$\Sigma_d : \mathcal{M} \rightarrow \mathcal{M}_d$$

together with induced constraint map

$$\Phi_d = \Phi \circ \Sigma_d^{-1}.$$

The domain-admissible manifold is

$$\mathcal{M}_d = \{s_d \in \mathcal{M}_d \mid \Phi_d(s_d) = 0\}.$$

Let Γ_d denote the admissible trajectory space over \mathcal{M}_d .

Define the lifting operator

$$L_d : \Gamma_d \rightarrow \Gamma$$

such that

$$\Sigma_d \circ L_d = \text{id}_{\Gamma_d}.$$

A domain is said to be faithful if

$$\Pi_d(\gamma_d) = \Pi(L_d(\gamma_d)) \quad \forall \gamma_d \in \Gamma_d,$$

where Π_d is the induced domain preference functional.

Define a domain instantiation as the tuple

$$(\mathcal{M}_d, \Gamma_d, \Phi_d, \Pi_d).$$

Two domains $d_1, d_2 \in \mathcal{D}$ are structurally equivalent if there exists a diffeomorphism

$$\varphi : \mathcal{M}_{d_1} \rightarrow \mathcal{M}_{d_2}$$

such that

$$\Phi_{d_2} \circ \varphi = \Phi_{d_1}, \quad \Pi_{d_2}(\varphi \circ \gamma) = \Pi_{d_1}(\gamma).$$

Let \mathcal{D}/\sim denote the quotient of domains under structural equivalence.

A macrostructure \mathcal{A} is domain-invariant if

$$\forall d \in \mathcal{D}, \quad \exists \mathcal{A}_d \subset \Gamma_d \quad \text{such that} \quad L_d(\mathcal{A}_d) \subset \mathcal{A}.$$

This establishes domain instantiation as constraint-preserving specialization rather than ad hoc model construction.

S A Functorial and Categorical Formulation

Define a category **Sim** whose objects are simulation systems

$$\mathcal{S} = (\mathcal{M}, \Gamma, \Phi, \Pi, S),$$

where \mathcal{M} is an admissible manifold, Γ its trajectory space, Φ a constraint map, Π a preference functional, and S an entropy functional.

A morphism

$$f : \mathcal{S}_1 \rightarrow \mathcal{S}_2$$

consists of a pair $(f_{\mathcal{M}}, f_{\Gamma})$ such that

$$f_{\mathcal{M}} : \mathcal{M}_1 \rightarrow \mathcal{M}_2, \quad f_{\Gamma} : \Gamma_1 \rightarrow \Gamma_2,$$

with the following compatibility conditions:

$$\Phi_2 \circ f_{\mathcal{M}} = \Phi_1, \quad \Pi_2 \circ f_{\Gamma} = \Pi_1, \quad S_2 \circ f_{\Gamma} = S_1.$$

Let **Dom** \subset **Sim** be the full subcategory of domain-instantiated systems.

Define a functor

$$\mathcal{F} : \mathbf{Dom} \rightarrow \mathbf{Pref}$$

to the category **Pref** whose objects are ordered sets (X, \preceq) and whose morphisms preserve order.

For each $\mathcal{S} \in \mathbf{Dom}$,

$$\mathcal{F}(\mathcal{S}) = (\Gamma/\sim_{\Pi}, \preceq),$$

where \sim_{Π} denotes preference equivalence.

Define a second functor

$$\mathcal{E} : \mathbf{Sim} \rightarrow \mathbf{Poset}$$

mapping each simulation to its entropy-ordered quotient space

$$\mathcal{E}(\mathcal{S}) = (\Gamma/\sim_S, \prec).$$

A simulation system is coherent if the diagram

$$\mathbf{Sim} @> \mathcal{F} >> \mathbf{Pref} @V \mathcal{E} VV @VV \mathcal{U} V \mathbf{Poset} @>> \iota > \mathbf{Pref}$$

commutes up to natural isomorphism, where ι is the forgetful functor.

Define the macrostructure functor

$$\mathcal{M}ac : \mathbf{Sim} \rightarrow \mathbf{Set}$$

by

$$\mathcal{M}ac(\mathcal{S}) = \mathfrak{M},$$

the set of macrostructural fixed points of \mathcal{S} .

A simulation civilization corresponds to a terminal object

$$\mathcal{S}^* \in \mathbf{Sim}$$

such that for any \mathcal{S} there exists a unique morphism

$$\mathcal{S} \rightarrow \mathcal{S}^*$$

preserving admissibility, entropy ordering, and preference equivalence.

Such a terminal object represents a maximally refined, resolution-independent, domain-complete simulation infrastructure in which all admissible futures factor through a common macrostructural selection space.

This categorical formulation unifies constraint geometry, trajectory dynamics, irreversibility, preference revelation, and domain specialization into a single abstract object, establishing simulation as a universal coordination medium independent of narrative representation.

T Sheaf-Theoretic Localization and Gluing of Futures

Let \mathcal{M} be an admissible manifold equipped with a topology τ generated by constraint-compatible open sets.

For each open set $U \subset \mathcal{M}$, define a local simulation system

$$\mathcal{S}(U) = (\mathcal{M}|_U, \Gamma(U), \Phi|_U, \Pi|_U, S|_U),$$

where $\Gamma(U)$ denotes admissible trajectories whose image lies entirely in U .

Define a presheaf

$$\mathcal{S} : \mathbf{Open}(\mathcal{M})^{op} \rightarrow \mathbf{Sim}$$

assigning to each open set its local simulation system and to each inclusion

$$V \subseteq U$$

a restriction morphism

$$\rho_{U,V} : \mathcal{S}(U) \rightarrow \mathcal{S}(V).$$

The presheaf \mathcal{S} is a sheaf if for any open cover $\{U_i\}$ of U and any compatible family of local trajectories

$$\{\gamma_i \in \Gamma(U_i)\}$$

satisfying

$$\rho_{U_i, U_{ij}}(\gamma_i) = \rho_{U_j, U_{ij}}(\gamma_j) \quad \text{on overlaps } U_{ij},$$

there exists a unique global trajectory

$$\gamma \in \Gamma(U)$$

such that

$$\rho_{U,U_i}(\gamma) = \gamma_i \quad \forall i.$$

Define the sheaf of macrostructures

$$\mathcal{Mac}(U) = \mathfrak{M}(U),$$

where $\mathfrak{M}(U)$ denotes the set of macrostructural fixed points of $\mathcal{S}(U)$.

A global macrostructure $\mathcal{A} \in \mathfrak{M}(\mathcal{M})$ exists if and only if the family

$$\{\mathcal{A}_i \in \mathfrak{M}(U_i)\}$$

is compatible under restriction on overlaps and admits a unique gluing.

Define an obstruction class

$$\omega \in H^1(\mathcal{M}, \mathcal{Mac})$$

measuring failure of local macrostructures to glue globally.

Vanishing of ω is a necessary and sufficient condition for coherent global futures to emerge from locally admissible dynamics.

This establishes that global civic futures are not primitive objects but sections of a sheaf of simulations, existing only when local constraint satisfactions are mutually compatible.

U Homotopy of Futures and Equivalence of Trajectories

Let Γ be the admissible trajectory space endowed with a topology induced by the metric d_Γ and constrained by admissibility and entropy monotonicity.

Define a homotopy between two trajectories $\gamma_0, \gamma_1 \in \Gamma$ as a continuous map

$$H : [0, 1] \times [0, T] \rightarrow \mathcal{M}$$

such that

$$H(0, t) = \gamma_0(t), \quad H(1, t) = \gamma_1(t),$$

and for all $s \in [0, 1]$, the path

$$\gamma_s(t) = H(s, t)$$

lies in Γ .

A homotopy is admissible if

$$\mathcal{C}(\gamma_s(t)) = 1 \quad \text{and} \quad \frac{d}{dt} S(\gamma_s|_{[0,t]}) \geq 0 \quad \forall s, t.$$

Two trajectories are homotopy-equivalent if there exists an admissible homotopy between them. Define the homotopy class of a trajectory as

$$[\gamma]_{\sim_H} = \{\gamma' \in \Gamma \mid \gamma' \sim_H \gamma\}.$$

Let $\Gamma_H = \Gamma / \sim_H$ denote the space of homotopy classes.

Define the catastrophic exclusion condition: a homotopy H is forbidden if

$$\exists s, t \text{ such that } H(s, t) \in \mathcal{X},$$

where \mathcal{X} is the catastrophic boundary set.

Homotopy classes are therefore defined relative to $\Gamma_{\text{safe}} = \Gamma \setminus \mathcal{X}$.
 Assume the preference functional Π is homotopy-invariant:

$$\Pi(\gamma_0) = \Pi(\gamma_1) \quad \text{if} \quad \gamma_0 \sim_H \gamma_1.$$

Then Π descends to a well-defined functional

$$\tilde{\Pi} : \Gamma_H \rightarrow \mathbb{R} \cup \{-\infty\}.$$

Define a future type as a connected component of Γ_H .

Let $\pi_0(\Gamma_{\text{safe}})$ denote the set of path-connected components under admissible homotopy.

Two macrostructures $\mathcal{A}_1, \mathcal{A}_2$ are future-equivalent if

$$\forall \gamma_1 \in \mathcal{A}_1, \exists \gamma_2 \in \mathcal{A}_2 \quad \text{such that} \quad [\gamma_1]_{\sim_H} = [\gamma_2]_{\sim_H}.$$

The set of inequivalent futures is therefore classified by

$$\pi_0(\Gamma_{\text{safe}}).$$

This homotopy classification establishes that distinct futures are separated not by preference magnitude alone, but by topological obstructions in trajectory space, identifying civilizational transitions as crossings between homotopy classes that cannot be achieved without violating existential constraints.

References

- [1] Barbour, J. (1999). *The End of Time*. Oxford University Press.
- [2] Penrose, R. (1979). Singularities and Time-Asymmetry. In S. W. Hawking and W. Israel (Eds.), *General Relativity: An Einstein Centenary Survey*. Cambridge University Press.
- [3] Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press.
- [4] Penrose, R. (2010). *Cycles of Time: An Extraordinary New View of the Universe*. Bodley Head.
- [5] Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
- [6] Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell.
- [7] Hayek, F. A. (1945). The Use of Knowledge in Society. *American Economic Review*, 35(4), 519–530.
- [8] Samuelson, P. A. (1938). A Note on the Pure Theory of Consumer's Behaviour. *Economica*, 5(17), 61–71.
- [9] North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.
- [10] Collins, H., and Evans, R. (2007). *Rethinking Expertise*. University of Chicago Press.
- [11] Prigogine, I. (1980). *From Being to Becoming: Time and Complexity in the Physical Sciences*. W. H. Freeman.
- [12] Scott, J. C. (1998). *Seeing Like a State*. Yale University Press.
- [13] Edwards, P. N. (1996). *The Closed World: Computers and the Politics of Discourse in Cold War America*. MIT Press.
- [14] Gerovitch, S. (2002). *From Newspeak to Cyberspeak: A History of Soviet Cybernetics*. MIT Press.
- [15] Rawls, J. (1993). *Political Liberalism*. Columbia University Press.
- [16] Habermas, J. (1996). *Between Facts and Norms*. MIT Press.
- [17] Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575–599.
- [18] DeLanda, M. (2006). *A New Philosophy of Society: Assemblage Theory and Social Complexity*. Continuum.
- [19] Maskin, E. (2007). Mechanism Design: How to Implement Social Goals. *American Economic Review*, 97(2), 567–576.
- [20] Winsberg, E. (2010). *Science in the Age of Computer Simulation*. University of Chicago Press.
- [21] Helbing, D. (2012). Social Self-Organization. In D. Helbing (Ed.), *Social Self-Organization*. Springer.