

# Substrate Independent Thinking Hypothesis (SITH): Life, Meaning, and the Autocatalytic Closure of Information

Flyxion

January 11, 2026

## Abstract

Recent work in theoretical biology and complex systems has advanced a substrate-independent conception of life grounded not in Earth-specific chemistry, but in universal principles of optimization, physical constraint, information processing, and hierarchical organization. In particular, research associated with the Santa Fe Institute has emphasized that life is best understood as a learning dynamic operating under energetic and informational limits, rather than as a privileged material configuration.

This paper extends that framework by proposing the *Substrate Independent Thinking Hypothesis* (SITH): any sufficiently closed, recursively self-maintaining bundle of information constitutes a living system in a formal organizational sense, independent of biological, cognitive, cultural, or technological substrate. Under this hypothesis, collections of documents, symbols, procedures, institutions, and built environments may qualify as living entities insofar as they preserve identity, regulate admissible transformations, suppress entropy, and propagate themselves across time.

The hypothesis is developed in stages. First, it situates SITH relative to existing theories of life, including autopoiesis and autocatalytic sets, clarifying where it extends and where it diverges from these frameworks. It then generalizes autocatalytic closure beyond chemistry to semantic systems, analyzes documents as semantic organisms, and introduces a universal semantic error threshold governing informational collapse. These claims are formalized using category theory, information geometry, and the Relativistic Scalar–Vector–Entropy Plenum (RSVP) field framework.

Finally, the paper examines material instantiations of semantic life. Furniture, architecture, and infrastructure are analyzed as operators that externalize agency, enforce semantic persistence, and stabilize collective intention. A detailed treatment of bounded political spaces—popularly remembered as “the room where it happened”—illustrates how physical constraints generate durable semantic outcomes. The resulting theory reframes culture, bureaucracy, and the built environment as living semantic organisms, with implications for ethics, governance, platform design, and the conservation or destruction of meaning-bearing systems.

## 1 From Universal Life to Universal Thinking

The problem of defining life in universal terms has long been constrained by its historical association with terrestrial biology. Classical definitions emphasize metabolism, cellular organization, or specific biochemical substrates, thereby encoding contingent features of Earth-based life into what is intended

to be a general concept. While such definitions are operationally effective within biology, they become increasingly inadequate when extended to artificial systems, cultural evolution, or the search for life beyond Earth.

Over the past several decades, this limitation has been progressively addressed by theoretical approaches that reconceive life as an organizational and dynamical phenomenon. Rather than asking what life is made of, these approaches ask what life does and under what constraints it persists. On this view, life is a process that optimizes information propagation under physical limits, subject to scaling laws, energetic costs, and error thresholds. The specific materials in which this process is instantiated—DNA, proteins, membranes—are treated as historically contingent solutions rather than defining features.

This shift has significant consequences. Once life is understood as a substrate-independent organizational pattern, it becomes possible to identify life-like dynamics in systems that are not conventionally biological. However, most existing formulations stop short of fully extending this logic to thinking itself. Cognition, meaning, and symbolic organization are often treated as derivative phenomena layered atop biological life, rather than as life-like processes governed by the same principles.

The central claim of this paper is that this asymmetry is unwarranted. The same substrate-independent logic that applies to life applies equally to thinking. Thinking, in the sense relevant here, is not a privileged activity of brains, nor does it require consciousness, subjectivity, or intentional experience. It is instead a class of informational dynamics characterized by recursive structure, regulated transformation, and resistance to entropic degradation. Wherever such dynamics occur, a form of life is present in a precise organizational sense.

This motivates the Substrate Independent Thinking Hypothesis. Informally stated, SITH holds that thinking systems are living systems whose primary mode of persistence is semantic rather than metabolic. Under this hypothesis, collections of documents, legal codes, rituals, technical standards, scientific theories, and architectural plans can satisfy the same viability conditions as biological organisms. The claim is not that such systems are conscious or sentient, but that they instantiate life at the level of informational organization.

Crucially, SITH does not assert that “everything is alive.” It asserts that life emerges wherever informational closure is achieved under constraint. Many symbolic systems fail to meet this standard and rapidly dissolve into noise. Others persist for centuries, adapt to changing environments, and exert causal influence across generations. The difference lies not in symbolism or belief, but in organizational closure.

The remainder of this paper develops this claim systematically. The next subsection situates SITH relative to existing theoretical frameworks in biology and cognition, clarifying its intellectual lineage and points of divergence. Subsequent sections generalize autocatalytic closure beyond chemistry, analyze documents as semantic organisms, introduce a universal semantic error threshold, and formalize these ideas using category theory, information geometry, and field dynamics. Only after this formal groundwork is established does the paper turn to material environments—furniture, rooms, buildings, and cities—which are treated not as metaphors but as operators that embed semantic viability into

physical space.

### 1.1 Relation to Existing Frameworks

The Substrate Independent Thinking Hypothesis does not arise in isolation. It builds directly on several established frameworks in theoretical biology and cognitive science while extending them into domains they were not originally designed to address. Clarifying these relationships is essential both for situating the present work and for distinguishing its claims from neighboring theories.

Autopoietic theory, as developed by Maturana and Varela, characterizes living systems as networks of processes that recursively produce and maintain the components that constitute the network itself. This emphasis on organizational closure rather than material composition aligns closely with SITH. However, classical formulations of autopoiesis remain tightly coupled to biochemical realization and offer limited guidance for analyzing symbolic or institutional systems. SITH generalizes autopoiesis by abstracting closure away from material production and toward the recursive maintenance of informational identity.

Stuart Kauffman’s theory of autocatalytic sets provides a second key precursor. Autocatalysis demonstrates how self-sustaining organization can emerge from sufficiently rich networks of interaction without centralized control. SITH adopts this insight but extends it from chemical reaction networks to semantic transformation networks. Where autocatalytic molecules catalyze reactions, semantic systems catalyze interpretations, procedures, and institutional actions. The underlying logic of closure is preserved, even as the substrate changes.

Recent work emphasizing scale, information, and thermodynamic efficiency in living systems further motivates the present approach. By treating life as an optimization process constrained by physics, these accounts open the door to genuinely substrate-independent formulations. SITH can be understood as an extension of this program into the domain of meaning, where the relevant constraints are informational rather than metabolic.

The principal point of divergence lies in scope. Existing frameworks aim to explain life as traditionally conceived, whereas SITH proposes that the same formal conditions govern a broader class of systems. This proposal is speculative, but it is not unconstrained. The sections that follow develop explicit mathematical criteria for semantic viability and show how these criteria can, in principle, be tested and falsified.

## 2 Autocatalytic Closure Beyond Chemistry

Autocatalytic closure has played a central role in contemporary theories of life, particularly in accounts of life’s origins. In its classical formulation, an autocatalytic system consists of a network of reactions in which the products of some reactions catalyze others, forming a closed loop that is collectively self-sustaining. No individual reaction is sufficient to constitute life; rather, life emerges only when the network as a whole achieves closure under its own dynamics.

The significance of autocatalysis lies not in chemistry per se, but in organization. What matters is not the material nature of the reactions, but the fact that the system generates the conditions of its

own persistence. Life, on this view, is not something added to matter; it is a pattern of constraint that allows matter to maintain itself against entropy.

The Substrate Independent Thinking Hypothesis adopts this insight while generalizing it beyond its chemical origins. Under SITH, autocatalysis is understood as a property of transformation systems more generally. A system is autocatalytically closed when its internal transformations collectively regenerate the conditions under which those transformations remain admissible. Once this abstraction is made, the domain of application expands dramatically. Any system of transformations—chemical, symbolic, procedural, or institutional—may in principle achieve closure.

In semantic systems, the role played by chemical reactions is taken up by meaning-preserving transformations. These include interpretation, translation, citation, amendment, compilation, execution, and enforcement. When a collection of symbolic elements is organized such that these transformations tend, in aggregate, to reproduce and stabilize the system’s own identity, a semantic autocatalytic system is formed. Closure, in this context, refers to the fact that the system’s internal rules of transformation are sufficient to maintain that identity over time.

This generalization does not eliminate the need for substrates. Symbols must be written, spoken, stored, executed, or embodied somewhere. However, the substrate functions as a carrier rather than as a defining feature. The same semantic system may migrate across substrates—oral, written, printed, digital—while preserving its organizational identity. Substrate independence applies to the level of organization, not to the level of physical implementation.

Failure, correspondingly, is also substrate-independent. Semantic systems collapse when their transformation networks cease to regenerate identity. This may occur through excessive variation, loss of correction mechanisms, or environmental perturbations that overwhelm internal constraints. Collapse is thus a loss of closure rather than a loss of material support.

This perspective provides the conceptual foundation for treating semantic systems as candidates for life under SITH. The next section develops this claim concretely by examining documents and document collections as paradigmatic examples of semantic autocatalytic organization.

### 3 Documents as Semantic Organisms

Documents are conventionally treated as passive containers of meaning, serving as vehicles through which information is transmitted from authors to readers. While this view captures an important aspect of their function, it obscures the dynamics by which documents persist, transform, and exert causal influence across time. Under SITH, documents are not merely representations; when embedded in sufficiently closed transformation networks, they function as components of living semantic organisms.

A single document, considered in isolation, rarely qualifies as a living system. Like an isolated molecule, it lacks the capacity to regenerate itself under transformation. However, when documents are embedded within structured ecologies of interpretation, citation, revision, and enforcement, they participate in networks that can maintain identity across generations of carriers. Constitutions, legal codes, religious canons, scientific theories, technical standards, and programming languages exemplify such systems. Their persistence does not depend on the continued existence of any particular individual,

but on the stability of the transformation pathways that reproduce them.

These systems exhibit behaviors that closely parallel those of biological organisms. They reproduce through copying, dissemination, citation, and institutional adoption. They vary through amendment, reinterpretation, translation, and contextual application. They compete for attention, legitimacy, and resources within shared environments. They maintain boundaries through mechanisms such as canonicity, accreditation, orthodoxy, version control, and procedural validity. None of these processes requires centralized intent or conscious design at the level of the system as a whole. They emerge from the interaction of local rules and constraints.

What distinguishes a living semantic system from a mere aggregation of texts is the regulation of admissible transformation. In a viable system, not all changes are permitted. Some transformations are recognized as preserving identity, while others are treated as errors, corruptions, or departures. This distinction may be enforced explicitly, as in legal or technical systems, or implicitly, as in linguistic or cultural practice. In either case, the regulation of transformation functions as an immune system, suppressing variations that would otherwise dissolve identity.

This analysis clarifies why document systems can exert causal power independently of individual belief. Laws constrain behavior not because all actors assent to them, but because institutional machinery continues to interpret, apply, and enforce them. Scientific paradigms guide research not because they are universally agreed upon, but because they structure training, publication, and evaluation. The agency of such systems resides not in minds alone, but in the persistence of organizational closure.

Treating documents as semantic organisms also reframes questions of preservation and decay. Archival practices, editorial standards, peer review, and institutional memory function as error-correction mechanisms that reduce effective mutation rates. Conversely, environments that privilege novelty, speed, or engagement over coherence increase variation and risk crossing critical thresholds beyond which identity can no longer be maintained. These dynamics are structural rather than moral, and they can be analyzed quantitatively once appropriate formal tools are introduced.

The next section introduces one such tool by generalizing the concept of the biological error threshold to semantic systems. This generalization provides a precise criterion for semantic collapse and prepares the ground for the formal treatments that follow.

## 4 The Error Threshold Reappears as Semantic Collapse

Any theory that treats semantic systems as living entities must account not only for their persistence, but also for their failure. In biological contexts, this role is played by the concept of the error threshold, originally developed in the study of molecular evolution. The error threshold captures a fundamental limitation: information-bearing systems can sustain only a finite amount of complexity given a nonzero rate of error in replication. Beyond a critical mutation rate, selection can no longer preserve the informational signal, and the system collapses into noise.

The Substrate Independent Thinking Hypothesis generalizes this principle to semantic systems. Documents, languages, institutions, and cultural practices all propagate information through processes that introduce variation. Copying, interpretation, translation, execution, and contextual application

are never perfectly faithful. Each transformation introduces the possibility of drift, ambiguity, or reinterpretation. When such variation is sufficiently constrained and corrected, semantic identity is preserved. When it is not, coherence degrades.

The semantic error threshold can therefore be understood as a boundary in transformation space. Below the threshold, internal correction mechanisms—such as precedent, ritual repetition, formal validation, or institutional enforcement—are sufficient to stabilize identity. Above it, uncontrolled variation accumulates faster than repair can compensate, and the system loses the capacity to reproduce itself as the same system. Importantly, this collapse does not require malicious intent, ideological corruption, or psychological failure. It is a structural consequence of the balance between informational load and transformation fidelity.

This framework reframes a wide range of familiar phenomena. Languages fragment when rates of borrowing and innovation exceed the capacity of shared norms to stabilize meaning. Legal systems decay when amendment, reinterpretation, or selective enforcement outpaces institutional correction. Scientific fields lose coherence when publication volume overwhelms review, replication, and synthesis. In each case, collapse is not a moral failure but a dynamical one. The system has exceeded its viable operating regime.

The semantic error threshold also clarifies the functional role of conservatism in semantic systems. Mechanisms that resist change—such as strict canonicity, slow amendment procedures, or rigid formatting requirements—often appear regressive or authoritarian when viewed normatively. Structurally, however, they function as error-correction strategies. By reducing effective mutation rates, they allow complex systems to persist under noise. Whether such strategies are adaptive or maladaptive depends on environmental conditions, not on their mere existence.

Crucially, the error threshold is substrate-independent. It applies to genetic sequences, neural representations, linguistic corpora, legal codes, and algorithmic systems alike. What differs across domains is not the existence of the threshold, but the location of the boundary and the mechanisms available for correction. This universality is a central claim of SITH and will be formalized explicitly in subsequent sections.

#### 4.1 Objection and Response: Are Semantic Systems Mind-Dependent?

A natural objection to the Substrate Independent Thinking Hypothesis is that semantic systems appear to require minds as their substrates. Words must be understood, documents must be interpreted, and symbols must be recognized as meaningful. From this perspective, semantic systems would seem to be parasitic on cognitive agents rather than independent entities. If this objection holds, the claim of substrate independence would be undermined.

The force of this objection rests on a conflation of implementation with organization. While semantic systems require physical instantiation and interaction with cognitive agents, the same is true of biological systems with respect to chemistry and physics. Cells require molecules, but life is not reducible to any particular molecule. The relevant question is not whether a system requires a substrate, but whether its defining properties depend on the specific nature of that substrate.

Under SITH, substrate independence applies at the level of organizational pattern rather than

material realization. A semantic system is defined by the structure of its transformation network and the constraints that regulate those transformations. Minds function as carriers, actuators, and interpreters within this network, but they do not define its identity. Legal systems persist despite turnover among judges and citizens; programming languages persist despite changes in hardware and user communities; religious canons persist despite variation in belief and interpretation. In each case, identity is maintained by institutional and procedural constraints rather than by any particular mental state.

Dependence, moreover, does not imply lack of autonomy. Many biological organisms depend on symbiotic partners or host environments while retaining distinct evolutionary trajectories. Viruses require host cells to replicate, yet they are not identical to those cells. Semantic systems similarly depend on minds for enactment while maintaining organizational closure that exceeds any individual participant.

This distinction becomes clearer as semantic processes are increasingly automated. Version control systems, continuous integration pipelines, algorithmic enforcement mechanisms, and machine-mediated communication routinely execute meaning-preserving transformations without direct human interpretation at each step. These developments do not create semantic life *ex nihilo*, but they make explicit that the dynamics of meaning are not exhausted by subjective experience.

The objection that semantic systems require minds therefore does not invalidate SITH. It instead clarifies the role of minds as one class of substrates among others. Substrate independence is a claim about the portability of organizational closure, not the absence of physical realization.

The next section examines the minimal conditions under which such closure can arise. In particular, it analyzes why semantic systems require more than a single carrier, and how triangulation among multiple constrained participants enables the emergence of self-correcting meaning.

## 5 “Wherever Two or Three Are Gathered”: Minimal Closure and Triangulation

If semantic life is characterized by autocatalytic closure under constrained transformation, then it must be possible to specify the minimal conditions under which such closure can arise. In biological contexts, this question appears as the problem of how many interacting components are required for a self-sustaining network to emerge. In semantic contexts, it appears as a question of how many carriers are required for meaning to stabilize.

A single inscription, however carefully preserved, cannot correct its own drift. Without comparison, variation cannot be distinguished from corruption. Two carriers introduce redundancy and feedback, but they remain insufficient to stabilize identity. Divergence between interpretations can be detected, but not adjudicated. Each carrier may attribute error to the other without recourse to a shared standard.

The introduction of a third constrained carrier marks a qualitative transition. With three participants bound by a shared symbolic attractor, triangulation becomes possible. Disagreement can be localized, majority consistency can be identified, and deviation can be corrected relative to a reference that no single carrier controls. At this point, semantic closure becomes self-sustaining. Meaning no

longer depends on the reliability of any individual participant; it is distributed across the network.

This transition has long been recognized culturally and religiously, often expressed in compressed symbolic form. The claim that presence emerges “wherever two or three are gathered in a name” can be read, in this context, as an observation about minimal conditions for semantic life. The “name” functions as an identity attractor, while the gathering functions as a redundancy structure that suppresses drift. The presence that emerges is not supernatural, but organizational: a self-correcting semantic system.

The importance of triangulation extends far beyond small groups. Institutions, disciplines, and traditions all rely on mechanisms that perform the same function at scale. Peer review, appellate courts, councils, and distributed consensus protocols formalize triangulation to stabilize meaning across populations. Where such mechanisms fail or are removed, semantic systems become fragile, prone to fragmentation or capture.

Minimal closure thus provides a bridge between abstract theory and material instantiation. Once meaning requires triangulation to persist, it becomes advantageous to externalize that triangulation into durable structures. This pressure drives the emergence of documents, records, archives, and eventually architectural spaces that embed redundancy and constraint into the environment itself. The next section examines how this process unfolds across substrates.

## 6 Ontogenic Hacking as Substrate Migration

Ontogenic hacking refers to the capacity of living systems to circumvent the limitations imposed by their current substrates by inventing new architectures for information storage, transformation, and transmission. In biological evolution, genetic inheritance constrains both the rate and fidelity of information propagation. Nervous systems emerged as a means of storing and manipulating information within a single lifetime, reducing dependence on genetic change. Language further externalized information, enabling transmission between individuals without reproduction. Documents extended this process across generations, allowing information to persist independently of any particular body.

Under the Substrate Independent Thinking Hypothesis, ontogenic hacking is not merely an acceleration of evolution, but a migration of organizational closure. Each new substrate alters the balance between complexity and fidelity by changing the effective mutation rate and the available mechanisms of error correction. The semantic error threshold is not eliminated by such migrations; it is displaced. What changes is the range of viable complexity that can be sustained.

This perspective clarifies why cultural and technological systems can evolve orders of magnitude faster than biological ones while remaining stable. Their stability does not arise from intrinsic robustness, but from the availability of externalized correction mechanisms. Writing, formal procedures, standardized formats, and automated validation all function as semantic immune systems. They suppress variations that would otherwise accumulate beyond control.

Ontogenic hacking also introduces new vulnerabilities. As semantic systems migrate into faster, more interconnected substrates, the potential for entropy injection increases. Rapid propagation can overwhelm correction mechanisms, pushing systems past their error thresholds even as their

expressive power grows. The same technologies that enable unprecedented preservation of meaning also enable unprecedented rates of distortion.

The pressure to externalize correction does not stop with documents. As semantic systems grow in scale and complexity, they increasingly rely on material environments to enforce constraints. Furniture, rooms, buildings, and cities emerge not as incidental contexts, but as adaptive responses to the problem of maintaining semantic closure under noise. These environments embed admissible transformations into physical form, reducing reliance on individual cognition.

The next sections formalize these dynamics mathematically. Category theory provides a language for describing semantic closure abstractly, information geometry quantifies identity and drift, and the RSVP field framework models semantic dynamics in space and time. Only after these formalisms are established does the paper return to material operators, treating furniture and architecture as necessary consequences of semantic life rather than illustrative metaphors.

## 7 Formalization of SITH Using Category Theory

The preceding sections motivate the Substrate Independent Thinking Hypothesis conceptually and historically. To render the hypothesis precise, it is necessary to specify the structural conditions under which a semantic system qualifies as living. Category theory provides a natural formal language for this task, as it abstracts away from internal constitution and focuses instead on relations, transformations, and closure under composition.

The guiding idea is that a semantic system is defined not by the symbols it contains, but by the transformations it admits while preserving identity. These transformations form a structured space whose closure properties determine whether the system persists under perturbation. Category theory allows this structure to be described without privileging any particular substrate or implementation.

Let  $(\text{Sem})$  denote a category whose objects are semantic states. A semantic state is understood as an inscription, configuration of inscriptions, or procedural specification together with a set of invariants that determine identity for that system. These invariants may include logical consistency, legal validity, syntactic well-formedness, procedural correctness, or institutional legitimacy, depending on the domain under consideration. Morphisms in  $(\text{Sem})$  are meaning-preserving transformations: translations, revisions, interpretations, executions, applications, or adjudications that respect the relevant invariants.

Composition of morphisms corresponds to sequential application of transformations. Identity morphisms correspond to transformations that leave a semantic state unchanged. The category  $(\text{Sem})$  thus encodes the admissible space of change for semantic systems.

Within  $(\text{Sem})$ , a candidate living semantic system is represented by a subcategory  $(M \subset \text{Sem})$  that is closed under composition and identity. Closure is essential. It ensures that applying admissible transformations to elements of the system yields results that remain within the system, up to semantic equivalence. This property formalizes the notion of organizational closure introduced earlier.

Replication is represented by an endofunctor  $[ R : M \rightarrow M, ]$  which maps each semantic object to a new instantiation of that object—such as a copy, enactment, or re-deployment—and maps each

morphism to the corresponding transformation between instantiations. The endofunctor ( $R$ ) need not preserve objects exactly; it may introduce controlled variation. What matters is that its image remains within ( $M$ ) up to semantic equivalence.

Environmental perturbations are represented by a functor  $[ E : M \rightarrow \text{Sem} ]$  which captures the effects of contextual change, reinterpretation, technological shifts, or adversarial pressure. In general,  $(E)$  may map objects outside of  $(M)$ . A semantic system is viable if the combined action of perturbation and replication restores it to  $(M)$  up to equivalence.

Formally, define the composite evolution operator  $[ F = R \circ E ]$ . The system represented by  $(M)$  is viable if there exists a nonempty subset  $(A \subset \text{Ob}(M))$  such that for every  $(X \in A)$ , there exists  $(Y \in A)$  with  $(F(X) \boxtimes Y)$ , where  $(\boxtimes)$  denotes semantic equivalence. In other words,  $(A)$  is an invariant set under the dynamics induced by  $(F)$ .

This condition constitutes a categorical formulation of life. Persistence is expressed as the existence of a fixed point or invariant subset under repeated cycles of perturbation and repair. Collapse corresponds to the absence of such an invariant set.

To illustrate this abstraction, consider the case of a constitutional system. The constitution itself, together with its recognized amendments and authoritative interpretations, constitutes an object in  $(\text{Sem})$ . Judicial rulings, legislative clarifications, and formal amendments constitute morphisms that preserve constitutional identity. Radical reinterpretations or extralegal actions fall outside the admissible morphism set.

The subcategory  $(M)$  consists of all semantic states recognized as valid instantiations of the constitution. Replication occurs when constitutional principles are cited, taught, or applied in new cases. Environmental perturbations arise from social change, political conflict, or technological innovation. The system remains alive so long as its internal mechanisms—courts, precedents, amendment procedures—are sufficient to absorb perturbations and restore identity. When these mechanisms fail, the constitutional system collapses or is replaced.

This example demonstrates how SITH captures the life of semantic systems without recourse to metaphor. Category theory provides a language for expressing semantic closure and viability in purely structural terms. However, this formulation does not yet quantify stability or specify how close a perturbed state remains to identity. To address these questions, a metric notion of semantic distance is required. The next section introduces information geometry as a means of quantifying identity, drift, and collapse in semantic systems.

## 8 Formalization of SITH Using Information Geometry

The categorical formulation of SITH specifies which transformations are admissible within a semantic system and how closure may be maintained under perturbation. However, it remains qualitative with respect to stability. To analyze how close a perturbed semantic state remains to identity, and to quantify the conditions under which collapse occurs, a metric structure is required. Information geometry provides such a structure by representing semantic systems as points on a statistical manifold equipped with a natural notion of distance.

A semantic system may be characterized operationally by the distributions of observable outcomes it generates. These outcomes vary by domain: linguistic systems generate word sequences, legal systems generate rulings and precedents, scientific fields generate publications and citation patterns, and institutions generate procedural actions. For a given semantic state, these observables define a probability distribution over a measurable space. Variation in meaning corresponds to variation in the parameters governing that distribution.

Let  $(\mathbb{X})$  denote a parameter manifold indexing a family of probability distributions  $(p(x|\mathbb{X}))$ , where  $(x)$  ranges over observable semantic outputs and  $(\mathbb{X} \in \mathbb{X})$  encodes the system's internal organization. The Fisher information metric on  $(\mathbb{X})$  is defined by  $[g_{ij}(\theta) = \mathbb{E}_{x \sim p(\cdot|\theta)} \left[ \frac{\partial \log p(x|\theta)}{\partial \theta_i} \frac{\partial \log p(x|\theta)}{\partial \theta_j} \right]]$ . This metric measures the distinguishability of semantic states.

Semantic identity is represented not by a single point in  $(\mathbb{X})$ , but by a compact region  $(B \subset \mathbb{X})$ , referred to as the identity basin. Variations within this basin are tolerated by the system's invariants, while variations outside correspond to semantic failure. Meaning-preserving transformations correspond to trajectories that remain within  $(B)$ .

To make this concrete, consider a simple two-parameter semantic system characterized by the parameter vector  $\theta = (\theta_1, \theta_2)$ , where  $\theta_1$  encodes lexical choice frequencies and  $\theta_2$  encodes syntactic structure frequencies within a standardized corpus. Suppose the observable distribution factorizes approximately as

$$p(x|\theta) = p_1(x|\theta_1) p_2(x|\theta_2).$$

Under this assumption, the Fisher information matrix is approximately diagonal, with entries

$$g_{11}(\theta) = \mathbb{E} \left[ \left( \frac{\partial \log p_1}{\partial \theta_1} \right)^2 \right], \quad g_{22}(\theta) = \mathbb{E} \left[ \left( \frac{\partial \log p_2}{\partial \theta_2} \right)^2 \right].$$

The infinitesimal Fisher distance between two nearby semantic states is then given by

$$ds^2 = g_{11}(\theta) d\theta_1^2 + g_{22}(\theta) d\theta_2^2.$$

This distance admits a direct empirical interpretation. If the value of  $ds^2$  is small, the two states generate nearly indistinguishable distributions of observable texts. If it is large, the difference is empirically detectable. Semantic drift can thus be quantified as accumulated Fisher distance along a trajectory of transformations.

Replication and interpretation introduce stochastic perturbations in parameter space. Each transformation induces a random displacement  $\eta$  with covariance matrix  $\Sigma$ , representing the effective mutation rate of the system. Repair mechanisms—such as editorial standards, procedural checks, or institutional enforcement—introduce deterministic flows that contract distances toward the interior of the identity basin. The expected evolution of semantic distance therefore reflects a balance between diffusion and contraction on the manifold.

The semantic error threshold introduced earlier can now be expressed geometrically. Let  $\lambda < 1$  denote the average contraction factor induced by repair mechanisms, and let  $\sigma^2 = \text{tr}(\Sigma)$  denote the effective noise strength per transformation. Viability requires that the expected Fisher distance from

a reference point in the identity basin  $\mathcal{B}$  decrease on average after each cycle of perturbation and repair. When  $\sigma^2$  exceeds a critical value determined by the curvature and size of  $\mathcal{B}$ , this condition fails. Trajectories exit the basin with high probability, and semantic identity collapses.

This formulation clarifies the role of redundancy and triangulation. Multiple independent carriers reduce effective noise by averaging perturbations, thereby shrinking  $\sigma^2$ . The result is an expansion of the viable identity basin and a lower probability of threshold crossing. The same mechanism underlies error correction in communication systems and robustness in biological populations.

Information geometry thus provides a quantitative bridge between abstract notions of semantic closure and observable behavior. While precise estimation of semantic manifolds remains challenging, the framework specifies what must be measured in principle. This makes SITH empirically accessible rather than purely philosophical.

The next section introduces a dynamical field-theoretic formulation that integrates coherence, propagation, and entropy into a unified model. This formulation, known as the Relativistic Scalar–Vector–Entropy Plenum, provides a means of studying semantic life as a spatiotemporal process.

## 9 Mapping SITH onto RSVP Fields

The categorical and information-geometric formalisms developed in the preceding sections specify the structural and metric conditions for semantic viability, but they remain largely static. To analyze how semantic systems evolve in time, respond to perturbations, and organize themselves spatially, it is useful to introduce a dynamical description in which coherence, propagation, and entropy are treated as interacting fields. The Relativistic Scalar–Vector–Entropy Plenum (RSVP) provides such a description.

Let  $(M)$  be a semantic manifold whose points correspond to local semantic states. The topology of  $(M)$  depends on the application. In linguistic systems,  $(M)$  may be approximated by a high-dimensional embedding space. In legal or institutional systems, it may be represented as a graph of admissible states endowed with a continuous relaxation. The assumption required here is not global smoothness, but local differentiability sufficient to define gradients and flows.

On the semantic manifold  $M$ , define three interacting fields. First, a scalar field  $\Phi(x, t)$  is interpreted as a semantic coherence potential. Regions of high  $\Phi$  correspond to states that strongly attract transformations back toward identity. Second, a vector field  $\vec{v}(x, t)$  is interpreted as a propagation flow, encoding how semantic content moves through carriers, channels, or institutional pathways. Third, a scalar field  $S(x, t)$  is interpreted as semantic entropy density, measuring local ambiguity, disorder, or uncontrolled variation.

These fields are not independent. Semantic coherence suppresses entropy; propagation can either reinforce coherence or generate disorder; entropy erodes coherence and destabilizes flow. A minimal RSVP dynamics that captures these interactions is given by the coupled system

$$\frac{\partial \Phi}{\partial t} = D_\Phi \nabla^2 \Phi - \alpha S + \beta |\vec{v}|^2 - \gamma \Phi,$$

$$\frac{\partial \vec{v}}{\partial t} = -\nabla \Phi - \delta \vec{v} + \kappa \nabla^2 \vec{v},$$

$$\frac{\partial S}{\partial t} = D_S \nabla^2 S + \mu |\nabla \vec{v}|^2 - \nu \Phi.$$

Here  $D_\Phi$  and  $D_S$  are diffusion coefficients, while  $\alpha, \beta, \gamma, \delta, \kappa, \mu$ , and  $\nu$  are positive constants governing coupling strengths and dissipation rates. The precise values of these parameters depend on the semantic domain under study, but their qualitative roles are universal.

The interpretation of these equations is straightforward. Semantic coherence diffuses locally, decays in the absence of maintenance, is degraded by entropy, and is reinforced by organized propagation. Propagation flow is driven down gradients of coherence potential, damped by frictional effects, and smoothed by diffusion. Entropy diffuses, is generated by turbulent or poorly regulated propagation, and is suppressed by coherent structure.

Boundary conditions encode architectural and institutional constraints. A tightly regulated semantic system, such as a legal code or a controlled archive, may be modeled with no-flux boundary conditions on the propagation field  $\vec{v}$  and the entropy field  $S$ , preventing uncontrolled inflow of entropy or outflow of propagation. Open systems, such as public discourse or social media platforms, permit entropy inflow and unregulated propagation at boundaries. These boundary conditions are not merely technical; they correspond directly to material and organizational features examined in later sections.

Semantic viability corresponds to the existence of attractor regions in the joint field space where the coherence potential  $\Phi$  remains positive and bounded, the propagation field  $\vec{v}$  remains structured rather than turbulent, and the entropy density  $S$  remains below a critical threshold. These attractors are the field-theoretic realization of the identity basins described in information-geometric terms. When entropy production overwhelms coherence generation, attractors disappear and semantic identity collapses. The RSVP equations thus provide a dynamical formulation of the semantic error threshold introduced earlier.

The RSVP framework also clarifies the role of scale. As the spatial extent of a semantic system increases, diffusion terms become more significant, and maintaining coherence requires either stronger coupling or more restrictive boundaries. This scaling behavior mirrors biological constraints on organism size and complexity. It also explains why large semantic systems—such as states, corporations, or empires—depend heavily on architecture, infrastructure, and procedural rigidity to remain viable.

The RSVP formulation makes explicit the continuity between abstract semantic dynamics and material instantiation. Coherence potential corresponds to what institutions call authority, legitimacy, or clarity. Propagation flow corresponds to communication, enforcement, and circulation. Entropy corresponds to ambiguity, inconsistency, and noise. These quantities are not metaphorical; they are dynamical variables whose interactions determine whether a semantic system lives or dies.

Before turning to architectural operators, it is useful to examine an intermediate-scale system in which these dynamics can be observed empirically. The next section analyzes version-controlled repositories as semantic systems governed by RSVP dynamics, bridging formal theory and material practice.

## 10 An Intermediate-Scale Worked Example: Version-Controlled Repositories as Semantic Systems

Before turning to the analysis of furniture, architecture, and bounded political space, it is useful to examine a system that occupies an intermediate scale between abstract semantic formalisms and material instantiation. Version-controlled repositories provide such a case. They are explicitly designed to manage semantic persistence under transformation, and their operational rules make visible many of the dynamics that remain implicit in cultural or institutional systems.

A repository consists of a structured collection of documents together with a formalized space of admissible transformations. Commits, merges, rebases, and rollbacks are not merely technical operations; they constitute a grammar of meaning-preserving change. Each operation is constrained by rules that determine whether the resulting state is considered valid. These rules encode a notion of semantic identity that is independent of any particular contributor.

From the categorical perspective developed earlier, a repository corresponds to a subcategory ( $M \subset Sem$ ). Objects are repository states, each consisting of a specific configuration of files together with metadata such as history and permissions. Morphisms are commits and merges that satisfy syntactic and procedural constraints. Identity morphisms correspond to null operations, while composition corresponds to sequential application of changes.

Replication occurs whenever the repository is cloned, forked, or deployed. These operations instantiate new copies of the semantic system while preserving its organizational identity. Environmental perturbations take the form of conflicting contributions, divergent branches, or external dependencies that introduce instability. Viability depends on the presence of repair mechanisms that can absorb these perturbations without destroying identity.

Information geometry makes this process quantitative. Each repository state induces a distribution over observable behaviors: test outcomes, build artifacts, performance metrics, or interface contracts. Commits introduce perturbations in parameter space, while continuous integration systems act as contraction mechanisms that pull the system back toward a viable basin. When testing and review are robust, the effective mutation rate remains below the semantic error threshold. When they are bypassed or overwhelmed, drift accelerates and the repository fragments.

The RSVP framework renders these dynamics spatially and temporally explicit. The coherence potential  $\Phi$  corresponds to architectural clarity, modularity, and test coverage. The propagation field  $\vec{v}$  corresponds to the rate and direction of contributions, merges, and deployments. Semantic entropy  $S$  corresponds to unresolved conflicts, undocumented behavior, and accumulating technical debt. Healthy repositories are characterized by stable regions of high coherence potential that channel propagation into controlled pathways while suppressing entropy. Failing repositories exhibit turbulent flows, diffuse coherence, and rising entropy density.

This example demonstrates that semantic life need not be speculative or metaphorical. Version-controlled repositories are engineered semantic organisms whose viability conditions are well understood by practitioners, even if not typically described in biological terms. The same principles govern more traditional semantic systems, but there they are often obscured by historical contingency and

cultural interpretation.

The relevance of this example to the remainder of the paper is twofold. First, it establishes that semantic systems can be analyzed as living entities without invoking subjective intention or anthropomorphic agency. Second, it shows that maintaining semantic viability at scale requires externalization of constraints into durable structures. In repositories, these structures are procedural and algorithmic. In other domains, they are physical.

The next section examines how such constraints are externalized into furniture. Far from being passive objects, furniture functions as a class of operators that shape semantic dynamics by embedding admissible transformations directly into material form.

## 11 Furniture as Operators in the Substrate Independent Thinking Hypothesis

The analysis developed thus far establishes that semantic systems persist by externalizing constraints on admissible transformation. Documents, procedures, and algorithmic pipelines perform this function abstractly. Furniture performs it materially. The claim that furniture participates in semantic life is therefore not rhetorical. It follows directly from the requirement that viable semantic systems must embed correction, partitioning, and boundary enforcement into their environments.

Furniture may be defined formally as a class of material operators that induce constraints on semantic state transitions. Unlike documents, which primarily encode rules, furniture enforces them by shaping the physical affordances through which semantic activity occurs. In doing so, furniture alters the effective topology of the semantic manifold and the boundary conditions of the RSVP fields defined earlier.

Consider first the case of storage furniture such as drawers, cabinets, folders, and shelves. These objects impose spatial partitioning on symbolic artifacts. By doing so, they reduce entropy locally by segregating states that would otherwise interfere. The act of placing an object in a drawer is not merely spatial; it is classificatory. It asserts that the object belongs to a category, that it is retrievable under certain conditions, and that its relations to other objects are constrained.

From the categorical perspective, furniture induces a restriction on admissible morphisms. Objects stored in separate compartments cannot be arbitrarily composed without an explicit act of retrieval and recombination. This restriction suppresses uncontrolled transformations, reducing the effective mutation rate of the semantic system. In information-geometric terms, furniture reshapes the identity basin by increasing its curvature at the boundaries, making drift more costly.

The archetypal example is the junk drawer. At first glance, it appears maximally disordered. Yet junk drawers are not entropy sinks; they are entropy buffers. They localize heterogeneity within a bounded region, preventing it from contaminating the rest of the system. Items placed in a junk drawer are not destroyed; they are retained in a liminal state, available for reuse under future, unforeseen transformations. The drawer thus functions as a reservoir of latent semantic potential.

In RSVP terms, the junk drawer corresponds to a region of elevated entropy density ( $S$ ) that is nonetheless bounded by strong coherence gradients ( $\nabla\mathbb{X}$ ). Entropy is permitted locally in order

to preserve global coherence. This is precisely the strategy employed by biological systems that compartmentalize waste or variability to maintain organism-level identity. The junk drawer is therefore alive in the precise sense that it participates in the autocatalytic closure of the household's semantic system.

More structured furniture, such as filing cabinets and bureaus, implement stronger constraints. A filing cabinet does not merely store documents; it enforces a hierarchy of categories. This hierarchy determines which transformations are easy, which are difficult, and which are effectively forbidden. Documents placed under one heading are more likely to be composed with each other than with documents stored elsewhere. Over time, this bias shapes the evolution of policy, practice, and interpretation.

In large-scale institutions, control over furniture is control over semantics. The individual who determines what goes in which drawer, who has access to which cabinet, and how retrieval is authorized effectively governs the space of admissible transformations. This is why bureaucracies are defined not primarily by ideology, but by filing systems. Authority is exercised through control of partitions.

This observation generalizes to cities and empires. Archives, ministries, data centers, and record rooms function as semantic organs. Their furniture defines the topology of administrative memory. When regimes change, revolutions often begin not with laws, but with the seizure or destruction of records. Such acts are lethal to semantic organisms because they destroy the material operators that sustain closure.

Furniture also functions dynamically. Desks orient bodies toward documents. Tables structure interaction. Seating arrangements encode hierarchies of attention and authority. These effects are not symbolic in a weak sense; they are causal. They bias propagation flows, represented by the vector field  $\vec{v}$ , by making some interactions effortless and others arduous. Over time, these biases accumulate into stable patterns of agency.

The relevance of furniture to SITH is therefore unavoidable. Once semantic life exceeds the scale at which individual cognition can maintain closure, it must externalize constraint into matter. Furniture is one of the earliest and most persistent forms of such externalization. It is the interface between symbolic order and physical space.

This analysis prepares the ground for a more specific and culturally salient case: the bounded room as a composite semantic operator. If furniture enforces local constraints, rooms enforce global ones. The next section examines how rooms function as higher-order closures, and why decisive historical events so often occur within them.

## 12 The Room Where It Happened: Rooms as Compositions of Semantic Operators

If furniture functions as a class of local semantic operators, then rooms function as higher-order closures that compose these operators into a unified constraint system. A room is not merely a container for furniture. It is a bounded semantic domain in which admissible transformations are collectively restricted, amplified, or suppressed. Under the Substrate Independent Thinking Hypothesis,

rooms constitute one of the simplest material forms capable of enforcing global semantic closure.

Formally, a room may be understood as the categorical product of its constituent furniture operators together with boundary conditions imposed by walls, entrances, and controlled apertures. Each piece of furniture constrains a subset of semantic transitions; the room composes these constraints into a joint admissibility structure. What emerges is not reducible to any individual operator. The room defines a space of possibility that exists only when the operators act together within a bounded environment.

This composition has immediate consequences for semantic dynamics. Once a system enters a room, certain transformations become easy, others difficult, and still others impossible. The boundaries of the room eliminate degrees of freedom. Entry and exit points regulate propagation flow. Internal arrangement channels attention and interaction. In RSVP terms, rooms impose sharp boundary conditions on the propagation field  $\vec{v}$  and the entropy field  $S$ , while shaping the spatial distribution of the coherence potential  $\Phi$ . Coherence is intensified not by persuasion, but by enclosure.

This analysis provides a precise interpretation of the cultural intuition captured in the phrase “the room where it happened.” The significance of such rooms does not derive from secrecy alone, nor from the moral weight of the participants. It derives from the fact that decisive semantic transformations require environments in which entropy can be suppressed long enough for irreversible commitments to be made. The room is the mechanism by which this suppression is achieved.

The reference to arrangement—of menu, venue, and seating—is analytically central. Seating arrangements encode asymmetries of access, attention, and authority. The order of speakers, the direction of gaze, and the proximity of bodies all bias the propagation field  $\vec{v}$ . These biases accumulate into gradients of coherence potential that favor certain outcomes over others. Decisions reached under such conditions are not merely agreed upon; they are stabilized.

From a categorical perspective, the room functions as a functor that maps unconstrained interaction spaces into a constrained subcategory of  $(\text{Sem})$ . Within this subcategory, certain morphisms—compromises, bargains, ratifications—exist that are not available outside it. Once enacted, these morphisms may be propagated outward, but they cannot be reproduced without reconstituting the same closure conditions. This asymmetry explains why some decisions are said to “only happen once.”

Information geometry further clarifies the irreversibility involved. The act of decision corresponds to a trajectory in semantic parameter space that crosses a narrow saddle region separating identity basins. Outside the room, stochastic perturbations would likely deflect the trajectory. Inside the room, constrained interaction reduces effective noise, allowing the system to traverse the saddle and settle into a new basin. Once there, return is improbable without comparable constraint.

Rooms thus operate as semantic phase-transition chambers. They are sites where semantic systems deliberately alter their own topology. Treaties, constitutions, mergers, and declarations are rarely produced in open, high-entropy environments. They require rooms precisely because rooms are devices for managing the semantic error threshold.

This understanding dissolves the apparent mystique surrounding political and institutional spaces. The power of rooms is not symbolic in a superficial sense. It is dynamical. Control over rooms is control over which semantic transitions can occur with sufficient fidelity to become binding. Exclusion

from the room is therefore not merely exclusion from discussion; it is exclusion from participation in semantic life at that scale.

The logic of rooms extends naturally to larger built environments. Buildings, campuses, cities, and infrastructures are composed of nested rooms and corridors that regulate movement, interaction, and memory. These environments externalize agency by embedding constraints into space. The next section generalizes this analysis, treating built environments as operators that extend collective intentionality across time and scale.

## 13 Built Environments as Agency-Extending Operators

The analysis of furniture and rooms establishes that semantic systems externalize constraints into material form in order to preserve identity under noise. Built environments represent the next stage in this process. A built environment is not merely a collection of rooms, nor a passive backdrop for action. It is an operator that extends agency across space, time, and population by embedding semantic constraints into durable structure.

Agency, in this context, should not be understood as subjective intention alone. Within the framework developed here, agency refers to the capacity of a system to bias future state transitions toward a restricted subset of possibilities. An agent is anything that reliably makes some futures more likely than others. Under this definition, built environments qualify as agents insofar as they shape trajectories of movement, interaction, and interpretation in systematic ways.

Formally, let  $(A)$  denote a space of possible actions available to a population prior to construction. A built environment induces an extension operator  $[ E : A \rightarrow A' ]$  where  $(A')$  is a constrained action space in which some actions are facilitated, others inhibited, and still others rendered impossible. This operator is not neutral. It reflects the intentions, values, and assumptions of its designers, but once instantiated it operates independently of them. The environment continues to execute those intentions long after individual participants have changed.

Homes provide a minimal illustration. The placement of a refrigerator externalizes nutritional planning. The stove externalizes thermal control. Storage spaces externalize memory. Corridors and doors regulate encounter rates. These features bias daily behavior without requiring deliberation. In RSVP terms, the home increases coherence potential by reducing entropy in routine action, while channeling propagation flows along predictable paths. The result is a stable semantic micro-organism that persists even as occupants change.

As scale increases, the same principles apply. Roads externalize coordination among strangers. Hospitals externalize care protocols. Schools externalize curricula. Courts externalize adjudication. Each institution embeds a specific subset of admissible transformations into physical and procedural form. Collectively, they constitute a distributed semantic organism whose agency exceeds that of any individual participant.

Cities emerge when these operators are composed at scale. A city is not defined by population alone, but by the topology of constraints that govern movement, communication, and memory. Zoning laws, transit systems, archival institutions, and bureaucratic layouts jointly determine which semantic

trajectories are viable. In the RSVP framework, cities correspond to extended regions of semantic manifold stabilized by dense coherence fields and regulated propagation flows. Entropy is not eliminated; it is managed through compartmentalization and hierarchy.

Empires represent the extreme case. At imperial scale, direct cognitive coordination becomes impossible. Semantic viability depends almost entirely on infrastructure. Archives, standards, and administrative architecture become organs. The loss of such organs is fatal. Empires collapse not only when defeated militarily, but when their semantic closure fails—when records are destroyed, procedures become inconsistent, or communication channels fragment. These failures correspond to crossings of the semantic error threshold at scale.

The continuity between homes, cities, and empires is not metaphorical. It is structural. Each level externalizes agency by embedding constraint into matter. Each level exhibits the same viability conditions: sufficient coherence potential, regulated propagation, and entropy suppression below critical thresholds. Each level can be analyzed using the same formal tools introduced earlier, differing only in parameter values and boundary conditions.

This perspective reframes the ethics of design and governance. Building is not merely construction; it is the creation of agents. Demolition is not merely destruction; it is semantic death. Policy decisions that alter infrastructure reshape the space of possible futures in ways that persist beyond individual intent. Recognizing built environments as living semantic systems therefore carries normative implications, but these implications follow from structure rather than ideology.

Having completed the theoretical and applied arc from abstract closure to material instantiation, the paper now turns to synthesis. The concluding section situates the Substrate Independent Thinking Hypothesis within the broader landscape of theories of life, meaning, and agency, and identifies directions for further formalization and empirical testing.

## 14 Conclusion: Life as Closure, Thinking as Survival

This paper has advanced the Substrate Independent Thinking Hypothesis as a generalization of substrate-independent life to the domain of meaning. Building on recent work in theoretical biology and complex systems, it has argued that life is not defined by its chemical constituents but by its organizational closure under constrained transformation. SITH extends this principle by demonstrating that semantic systems—documents, institutions, infrastructures, and built environments—can satisfy the same criteria of viability, persistence, and collapse.

The argument has proceeded by stages. Autocatalytic closure was generalized beyond chemistry to encompass semantic transformation networks. Documents were shown to function as semantic organisms when embedded in correction-regulated ecologies. The biological error threshold was reinterpreted as a semantic limit on complexity under noise. Minimal closure was formalized through triangulation, clarifying why semantic life requires distributed constraint rather than isolated cognition. Ontogenetic hacking was reframed as the migration of organizational closure across substrates, culminating in material externalization.

These conceptual claims were then formalized using category theory, information geometry, and

the RSVP field framework. Category theory captured closure, replication, and viability abstractly. Information geometry provided a metric notion of identity and drift. RSVP integrated coherence, propagation, and entropy into a dynamical model capable of spanning scales from repositories to empires. Furniture, rooms, and built environments emerged not as metaphors but as necessary operators for maintaining semantic life under increasing scale and noise.

The result is a reframing of culture, bureaucracy, and architecture as living semantic systems. Meaning persists not because it is believed, but because it is structurally protected. Agency extends not only through minds, but through constraints embedded in matter. Survival, in this expanded sense, is the capacity to maintain closure against entropy across generations of carriers.

This reframing has consequences. It dissolves sharp distinctions between biology and culture, between life and institution, between organism and environment. It suggests that destruction of archives, erosion of procedures, and degradation of infrastructure are not merely social harms but forms of semantic extinction. Conversely, preservation, standardization, and architectural care function as conservation biology for meaning-bearing systems.

At the same time, the theory does not collapse into mysticism or metaphor. At each stage, claims have been grounded in formal structure and dynamical constraint. Where speculation remains, it is bounded by explicit assumptions and testable implications. The ambition of SITH lies not in asserting that everything is alive, but in specifying precisely what conditions must be met for anything to be so.

### 14.1 Falsifiability and Empirical Engagement

A theory that claims broad applicability must confront the question of empirical testability. The Substrate Independent Thinking Hypothesis makes several claims that distinguish it from conventional accounts of culture and cognition, and these claims admit, at least in principle, of empirical evaluation.

First, SITH predicts the existence of sharp viability thresholds in semantic systems analogous to biological error thresholds. As mutation rates increase—whether through rapid reinterpretation, uncontrolled remixing, or accelerated propagation—systems should exhibit sudden rather than gradual loss of identity. This prediction can be tested by analyzing historical corpora, legal regimes, or software ecosystems for phase-transition-like behavior in coherence metrics.

Second, the theory predicts that redundancy and triangulation measurably reduce semantic drift. Systems with multiple independent correction pathways should exhibit lower effective mutation rates than systems relying on single authoritative channels. Comparative studies of institutions, languages, or repositories with differing governance structures could test this claim.

Third, RSVP dynamics predict that architectural and infrastructural interventions alter semantic trajectories independently of individual intent. Changes in layout, access control, or boundary permeability should produce observable shifts in propagation flow and entropy density. Such effects are already studied implicitly in urban planning and organizational theory; SITH provides a unified interpretive framework for them.

Fourth, the theory predicts that semantic collapse can occur even in the absence of adversarial pressure, purely as a consequence of scale and noise. This distinguishes SITH from moralized accounts

of institutional decay and suggests quantitative early-warning indicators based on rising entropy density or declining coherence potential.

Evidence against SITH would include the absence of identifiable thresholds, the failure of redundancy to reduce drift, or the persistence of semantic identity under arbitrarily high noise without increased constraint. The framework therefore exposes itself to refutation rather than insulation.

The empirical challenges are substantial. Measuring semantic manifolds, estimating Fisher metrics, and parameterizing RSVP fields remain open problems. Nevertheless, the framework specifies what must be measured, even if current tools are inadequate. This specificity marks a step beyond metaphor toward theory.

The paper concludes with an invitation rather than a closure. If thinking is a form of life, then the preservation, design, and governance of semantic systems become matters of survival. The task ahead is to formalize this insight further, to test it where possible, and to decide—collectively—which semantic organisms we are willing to sustain.

## A Mathematical Appendix

### A.1 Fixed-Point Viability and Existence of Semantic Closure

The main text characterizes semantic life as the existence of an invariant set under combined perturbation and repair. This section provides a constructive justification for that claim.

Let  $\mathcal{M} \subset \mathbf{Sem}$  be a subcategory representing a semantic system, and let  $\mathcal{F} = \mathcal{R} \circ \mathcal{E}$  be the composite evolution operator defined in Section 7. We seek conditions under which there exists a nonempty subset  $\mathcal{A} \subset \text{Ob}(\mathcal{M})$  such that  $\mathcal{F}(\mathcal{A}) \subseteq \mathcal{A}$  up to semantic equivalence.

Assume that  $\mathcal{M}$  admits a compact identity basin when embedded into a metric space via the information-geometric construction of Section 8. Specifically, let  $(\Theta, d_F)$  be the Fisher metric space associated with the semantic system, and let  $\mathcal{B} \subset \Theta$  be a compact, convex region representing admissible identity variation.

Suppose that environmental perturbations act as a stochastic map

$$E : \Theta \rightarrow \Theta$$

with bounded variance  $\sigma^2$ , and that repair acts as a deterministic contraction

$$R : \Theta \rightarrow \Theta$$

with Lipschitz constant  $\lambda < 1$  on  $\mathcal{B}$ . Then the expected displacement under one cycle of evolution satisfies

$$\mathbb{E}[d_F(R(E(\theta)), \theta_0)] \leq \lambda \mathbb{E}[d_F(E(\theta), \theta_0)],$$

for some reference point  $\theta_0 \in \mathcal{B}$ .

By standard results on stochastic contractions in compact metric spaces, there exists an invariant measure supported on  $\mathcal{B}$ . In categorical terms, this corresponds to the existence of a nonempty invariant subset  $\mathcal{A}$  under  $\mathcal{F}$ . Semantic life, as defined in the main text, therefore exists whenever perturbations are sufficiently bounded and repair sufficiently contractive.

Collapse occurs when these conditions fail, either because  $\sigma^2$  becomes unbounded or because  $\lambda \geq 1$ , eliminating contraction. This formalizes the viability criterion invoked throughout the paper.

### A.2 Bounds on the Semantic Error Threshold

The semantic error threshold introduced in Section 4 can be derived explicitly from the information-geometric formulation. Let  $\mathcal{B}$  be an identity basin of radius  $r$  in Fisher distance. Assume perturbations are isotropic with covariance  $\Sigma = \sigma^2 I$ , and that repair induces contraction by a factor  $\lambda$  per cycle.

After one cycle, the expected squared distance from the basin center satisfies

$$\mathbb{E}[d_F^2] \approx \lambda^2 d_F^2 + \sigma^2.$$

Viability requires that this quantity remain less than  $r^2$  in expectation. Solving for  $\sigma^2$  yields the bound

$$\sigma^2 < (1 - \lambda^2) r^2.$$

This inequality defines the semantic error threshold. When the effective mutation rate  $\sigma^2$  exceeds this bound, trajectories escape the identity basin with high probability, and semantic collapse ensues.

The dependence on basin radius clarifies why more richly structured systems can tolerate higher absolute noise while remaining viable. It also explains why rapid scaling without proportional increases in correction capacity is destabilizing.

### A.3 Derivation of the RSVP Free-Energy Functional

The RSVP dynamics introduced in Section 9 admit a variational interpretation. Define the semantic free-energy functional

$$\mathcal{F}[\Phi, \vec{v}, S] = \int_M \left( \frac{1}{2} |\nabla \Phi|^2 + \frac{1}{2} |\vec{v}|^2 + \alpha S - \beta \Phi \right) d\mu,$$

where  $d\mu$  is the volume measure on the semantic manifold  $M$ , and  $\alpha, \beta > 0$  are coupling constants.

The gradient flow of this functional under appropriate dissipation terms yields the qualitative form of the RSVP equations. The free energy decreases monotonically along trajectories corresponding to stable semantic evolution. Entropy production corresponds to local increases in  $\mathcal{F}$ , while coherence generation corresponds to decreases.

This formulation connects RSVP to existing variational principles in nonequilibrium thermodynamics and active inference, while remaining agnostic about cognitive implementation.

### A.4 Stability and Scaling of RSVP Dynamics

Linearizing the RSVP equations around a homogeneous fixed point  $(\Phi_0, \vec{v}_0 = 0, S_0)$  yields stability conditions of the form

$$\alpha S_0 < \gamma \Phi_0,$$

ensuring that coherence decay is dominated by repair rather than entropy production. Diffusion terms introduce scale dependence: as system size increases, maintaining stability requires either increased coupling strength or stricter boundary conditions.

This analysis explains why large semantic systems depend on rigid infrastructure and compartmentalization. Without such constraints, diffusion-driven entropy overwhelms coherence at scale.

## References

- [1] C. P. Kempes, J. A. West, L. M. A. Bettencourt, and G. B. West. The scaling of information in biological systems. *Proceedings of the National Academy of Sciences*, 118(13), 2021.
- [2] H. R. Maturana and F. J. Varela. *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel Publishing Company, 1980.
- [3] S. A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- [4] M. Eigen. Selforganization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, 58:465–523, 1971.
- [5] D. C. Dennett. *Darwin’s Dangerous Idea*. Simon & Schuster, 1995.
- [6] D. R. Hofstadter. *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, 1979.
- [7] L. Floridi. *The Philosophy of Information*. Oxford University Press, 2011.
- [8] S. I. Walker and P. C. W. Davies. The algorithmic origins of life. *Journal of the Royal Society Interface*, 10(79), 2013.
- [9] A. Kolchinsky and D. H. Wolpert. Semantic information, autonomous agency, and nonequilibrium statistical physics. *Interface Focus*, 8(6), 2018.
- [10] K. Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11:127–138, 2010.
- [11] M. J. D. Ramstead, K. J. Friston, and I. Hipólito. Is the free-energy principle a formal theory of semantics? *From Brains to Social Sciences*, Springer, 2018.
- [12] A. Mesoudi. *Cultural Evolution: How Darwinian Theory Can Explain Human Culture*. University of Chicago Press, 2011.
- [13] J. Henrich. *The Secret of Our Success*. Princeton University Press, 2015.
- [14] L.-M. Miranda. *Hamilton: An American Musical*. Public Theatre, 2015.