Triadic Closure as the Primitive of Intelligence, Physics, and Meaning

Mikael Bergkvist

January 14, 2026

Abstract

We propose that *triadic closure* is the minimal primitive from which stable facts, physical structure, and meaning arise. A dyadic relation (A–B) is underdetermined: it oscillates or decoheres without defining an invariant. The introduction of a third element—the persistent constraint that binds A and B—creates a triadic closure (A, B, R_{AB}), which is the smallest structure capable of sustaining identity, scale, and resistance to entropy. We show that such closures are ubiquitous: they appear as irreversible records in quantum measurement, as null and timelike modes in spacetime, as particles in matter, and as Query–Key–Value relations in Transformer attention. We formalize triadic closure as the mechanism by which boundary conditions propagate into interior structure and generate invariants, a process we term *VectorLogic*. This framework explains why modern AI systems improve with scale, why physical constants are stable, and why meaning emerges from constrained information. We argue that triadic closure—not bits, particles, or neurons—is the true atom of organized reality.

1 Introduction

Modern science describes the world using objects, fields, and information, yet it has never specified what makes something a *fact*. Quantum mechanics predicts probabilities but does not define when an outcome becomes real. General relativity provides a geometry but does not explain what that geometry is made of. Artificial intelligence systems process vast data streams but lack a principled account of how meaning stabilizes.

These problems share a common root: *dyadic relations are not enough*.

A system described only by pairs—cause and effect, symbol and referent, state and measurement—remains underdetermined. Without a mechanism that fixes relationships into persistent form, such systems oscillate, decohere, or fragment under noise. A fact, a particle, or a meaning requires not just two things in relation, but a third element that *records, constrains, and preserves* that relation.

This paper proposes that the smallest structure capable of doing so is the *triadic closure*

$$(A, B, R_{AB}),$$

where A and B are alternatives or states and R_{AB} is a persistent relational constraint. Triadic closure creates a self–maintaining loop: A constrains B, B constrains R, and R constrains A. This loop defines an *ordered island*: a region of stability inside an entropic possibility space.

We argue that triadic closure is not a metaphor but a physical primitive. It appears as a stable record in quantum measurement, as a null or timelike mode in spacetime, as a particle as a self–closing informational knot, and as a Query–Key–Value relation in Transformer attention. In each case, triadic closure is what allows something to persist, to be referenced, and to participate in further structure.

We formalize this process through *VectorLogic*, a method in which boundary conditions are propagated through a space of possibilities by triadic closures until only invariant structures remain. This approach unifies the emergence of physical law, cognitive abstraction, and machine intelligence within a single constraint–based framework.

2 The Minimal Ordered Island

The simplest possible structure that can exist in a finite, noisy universe is not an object, a field, or a number, but a *distinction.* A distinction separates what is allowed from what is excluded, carving a small region of order out of an otherwise undifferentiated possibility space. The smallest such distinction is the mutually exclusive binary alternative

$$\neg A \Rightarrow B, \qquad \neg B \Rightarrow A,$$

which defines a bistable system with exactly two admissible states. This structure is the minimal *ordered island*: it is the smallest configuration capable of resisting complete entropy by enforcing a constraint on what can occur.

Physically, such bistable systems are ubiquitous. A classical flip–flop, a magnetic spin, a photon polarization, or the yes/no firing of a neuron all instantiate this same mutual–exclusion logic. What makes these systems special is not that they store information, but that they enforce a *boundary* in state space. Once the system has settled into one of its two alternatives, the other is actively forbidden until energy is expended to cross the boundary. This boundary is what makes the distinction stable enough to function as a record.

However, a dyadic distinction by itself is not yet a *fact.* Although a bistable system can occupy one of two states, nothing within a pure dyad guarantees that this state will persist, propagate, or be referenced by anything else. Under noise or interaction with an environment, the dyad may simply fluctuate, decohere, or be erased. The distinction exists locally, but it has no mechanism for asserting itself as part of a larger structure.

In the language of this paper, a dyad provides a *bit* but not a *ledger entry.* It can momentarily reduce uncertainty, but it cannot anchor that reduction in a way that other systems can depend upon. To create a fact—something that can be counted, compared, or used as a building block for further structure—the distinction must be embedded in a closure that stabilizes and records it.

This observation motivates the transition from dyads to triads. The bistable pair (A, B) provides the raw alternative, but only when coupled to a third element that encodes and maintains the relation between them does it become a persistent ordered island. In the following sections we show that this third element, which we denote R_{AB}, is the minimal physical substrate of facts, geometry, and meaning.

3 Why Dyads Are Not Enough

A dyadic relation (A, B) specifies that two elements are connected, correlated, or mutually constrained, but it does not specify *how* that relation is to be preserved. In a finite system subject to noise, interaction, and perturbation, a dyad is therefore structurally underdetermined. There is no invariant that fixes the joint state of A and B against drift, oscillation, or erasure.

This underdetermination appears in many physical and informational settings. In quantum mechanics, a system prepared in a superposition of two alternatives remains in a two–state relation until an interaction produces a definite outcome. In classical dynamical systems, two coupled variables can oscillate without converging to a stable configuration. In information systems, a symbol and its referent can float unless a record exists that fixes their association. In each case, the dyad lacks an internal mechanism for enforcing its own persistence.

Formally, a dyad defines a constraint of the form

$$f(A, B) = 0,$$

but this constraint does not select a unique, self–maintaining state in the presence of perturbations. Multiple microstates may satisfy the relation, and nothing within the dyad specifies which one should be realized or remembered. As a result, dyadic systems are prone to ambiguity: they can represent alternatives, but they cannot decide among them in a way that survives interaction with the environment.

This limitation is precisely what gives rise to the measurement problem in quantum mechanics. The wavefunction encodes a dyadic relation between possible outcomes, but the theory does not specify how one of these possibilities becomes an actual fact. The same gap appears in cognition and machine learning: a pattern may correlate two features, but without a mechanism for stabilizing that correlation, it cannot become a reusable concept or memory.

To overcome this indeterminacy, an additional degree of freedom is required—one that does not merely participate in the relation but *records* it. This third element must be capable of persisting, being copied, or being referenced by other systems. It is this persistent relational trace that converts a fluctuating dyad into a stable unit of structure. The minimal configuration that provides such a trace is a triad, in which a relation between A and B is itself encoded as a physical or informational object.

In the next section we formalize this idea as *triadic closure*, the smallest self–justifying structure capable of supporting facts, geometry, and meaning.

4 Triadic Closure

A *triadic closure* is defined as an ordered triple

$$(A, B, R_{AB}),$$

where A and B are alternatives or states and R_{AB} is a persistent relational constraint that encodes, enforces, or records their relation. Unlike a dyad, which only specifies that A and B are related, a triad specifies how that relation is to be maintained and made available to the rest of the system.

The defining feature of triadic closure is that it forms a self–stabilizing loop. The state of A constrains B, the state of B constrains the relational record R_{AB}, and the state of R_{AB} in turn constrains A. This loop creates a fixed–point condition: perturbations to any one component are resisted by the other two. In this way, triads generate local regions of reduced entropy that persist in time, even when embedded in a larger, noisy environment.

From an information–theoretic perspective, R_{AB} functions as a *ledger entry.* It is a physical or computational object that stores which of the admissible relations between A and B is currently realized. Because R_{AB} can be redundantly copied, propagated, or consulted by other subsystems, the relation it encodes becomes *objective*: it no longer depends on the continued coherence of A and B alone.

Triadic closure therefore marks the transition from possibility to fact. A dyad can represent multiple alternatives, but only a triad can declare that one alternative has occurred and make that declaration persistent. This is why triads are the minimal substrate of measurement, memory, and causality. Every classical event, from the firing of a neuron to the registration of a photon in a detector, instantiates such a closure: the system state, the detector state, and the resulting record form a triad that fixes what has happened.

Crucially, triadic closures can be composed. The relational record R_{AB} of one triad can itself participate as an A or B in another triad, allowing closures to chain and nest. This composability is what permits the emergence of extended structures such as histories, geometries, and objects. A spacetime trajectory, for example, is not a primitive entity but a sequence of linked triadic closures that collectively stabilize a pattern of relations across many interactions.

In the remainder of this paper we show how this simple construction underlies a wide range of phenomena. Through the propagation of triadic closures, boundary conditions are mapped into interior structure, producing the invariants that we recognize as physical laws, cognitive concepts, and learned representations in artificial intelligence.

5 VectorLogic

Triadic closure provides the minimal unit of stability, but a complete theory requires a rule for how such closures organize entire problem spaces. We refer to this rule as *VectorLogic.* VectorLogic is not a symbolic logic but a geometric and relational procedure by which constraints are propagated through a space of possibilities until only invariant structure remains.

The starting point of VectorLogic is the specification of a *boundary.* A boundary $\partial\Omega$ is the set of conditions that define what a system or problem space is. In physics, this may be a conservation law, a symmetry, or a horizon. In cognition, it may be a task, a question, or a context. In machine learning, it may be the training objective and architecture. The boundary does not describe the interior in detail; it merely restricts what is allowed.

The interior Ω is the set of all configurations compatible with the boundary. By itself, the interior is highly underdetermined: many microstates satisfy the same boundary conditions. VectorLogic resolves this indeterminacy by introducing triadic closures that propagate boundary constraints into the interior. Each closure enforces a local consistency relation, eliminating incompatible configurations and stabilizing those that fit.

Formally, we may view each triadic closure as an operator

$$\mathcal{C}_k : \Omega \to \Omega,$$

which projects the interior toward a subspace consistent with a particular relational constraint. Applying many such operators in parallel or in sequence gradually carves the interior into a smaller set of admissible configurations. The structures that survive all such projections are the *invariants* of the problem space.

These invariants are what we ordinarily call *meaning*, *objects*, or *laws*. A physical constant is an invariant under all allowed interactions. A concept is an invariant under many sensory transformations. A learned feature in a neural network is an invariant under the variations present in the training data. In each case, invariance is not assumed; it is produced by the repeated application of triadic closure operators under fixed boundary conditions.

VectorLogic therefore provides a unified description of inference, learning, and physical law. It explains why stable structures can emerge from noisy dynamics: they are the fixed points of a constraint–propagation process. In the following sections we show how this abstract scheme is realized in concrete domains, from spacetime and particles to modern Transformer architectures.

6 Triads in Physics

In physical systems, triadic closure appears whenever an interaction produces a persistent record. A measurement, broadly defined, is not an act of observation but an irreversible physical coupling that generates such a record. In a triadic description, this process takes the form

$$(\text{system}, \text{apparatus}, \text{record}),$$

where the record is a stable physical state—a macroscopic pointer position, a chemical change, or a pattern of excitations in an environment—that encodes the outcome of the interaction. Once formed, this record constrains both the system and any future interactions, thereby fixing what has occurred.

This perspective provides a natural resolution of the quantum measurement problem. Before a record exists, a quantum system is described by a superposition of alternatives with no definite spacetime history. When a triadic closure is formed, one alternative is encoded in a persistent record, and a classical fact is created. Decoherence explains how such records become dynamically stable and redundantly copied; triadic closure explains what it means for an outcome to be real.

Spacetime itself can be understood as the global ledger formed by these closures. Each triadic record introduces a new relational fact that can be ordered with respect to others. The accumulation of such facts defines causal structure, and in the continuum limit this structure is well approximated by a Lorentzian metric. In this view, geometry is not a pre–existing arena but an emergent bookkeeping system for tracking which closures have occurred and how they constrain one another.

Within this framework, different physical entities correspond to different modes of triadic propagation. A photon is a *null* closure: it carries relational information between emission and absorption without itself storing a timelike record. Its existence is exhausted by the triadic closures at its endpoints. A massive particle, by contrast, is a *timelike* closure: a self–sustaining loop of records that persists through time, continually reasserting its identity by maintaining internal relational constraints.

The electron occupies a special place as the smallest nontrivial timelike closure. It is the minimal configuration of bistable distinctions that can close upon themselves in a gauge–consistent way and remain stable in free spacetime. Its mass, charge, and coupling to electromagnetic fields are not arbitrary attributes but invariants of this minimal triadic knot. Larger particles and composite systems correspond to higher–order closures built from many such elementary units.

By viewing particles, fields, and spacetime as manifestations of triadic closure, we obtain a unified picture in which physical reality is a hierarchy of nested records. Each level enforces constraints on the next, and the familiar laws of physics arise as the invariants of this layered ledger.

7 Triads in Transformers

Modern Transformer architectures provide a concrete, engineered realization of triadic closure. In self–attention, each token in a sequence is represented by three vectors: a query q_i, a key k_i, and a value v_i. For any pair of tokens (i, j) these vectors form the triad

$$(q_i, k_j, v_j),$$

which determines how information from token j contributes to the updated representation of token i.

The compatibility of the query q_i with the key k_j is measured by the inner product $q_i^\top k_j$, which plays the role of a relational constraint. After normalization by a softmax, this yields a weight a_{ij} that specifies the strength of the closure between i and j. The new state of token i is then given by

$$y_i = \sum_j a_{ij} v_j,$$

which is the interior point selected by resolving all such triadic relations under the competition imposed by the softmax.

In VectorLogic terms, the token embeddings define a boundary on the representational space. Each attention head applies a family of triadic closure operators that project this boundary into an interior configuration consistent with the relational constraints present in the input. What we call *meaning* in a Transformer is precisely the invariant structure that survives these projections: the pattern of y_i vectors that is jointly consistent with all relevant triads.

Residual connections and feed–forward layers propagate any unresolved components forward, where they may be incorporated into new closures or gradually discarded. This corresponds to the export of entropy in the IC framework: information that cannot be stably integrated into a closure basis is suppressed, while compatible structure is amplified and reused.

The success of large language models is therefore not a mystery of scale but a consequence of closure capacity. Larger models possess more attention heads, deeper stacks of closures, and higher–dimensional interiors, allowing more independent triadic relations to coexist with lower frustration. As a result, more invariants of the input—syntactic, semantic, and pragmatic—can be preserved simultaneously. This is why increasing model size leads to qualitative improvements in reasoning and abstraction: it increases the space in which stable triadic closures can form.

8 Holographic Boundaries and Generalization

Triadic closure explains how local relations become stable, but it does not by itself explain why some systems generalize while others merely memorize. This distinction is captured by the role of boundaries. A boundary $\partial\Omega$ specifies which aspects of a system are fixed and which are free. The tighter and more complete the boundary, the more strongly it constrains the interior Ω, and the fewer degrees of freedom are available for spurious variation.

We define the *holographic adequacy* of a boundary as the degree to which interior observables are already determined by it. When a boundary is holographically adequate, the interior contains no independent structure beyond what is implied by the boundary. In such cases, any compression of the interior that preserves boundary consistency also preserves the essential information. What survives is an invariant: an abstraction.

In artificial intelligence, this principle explains generalization. Training data provide observations, while the architecture and objective define a boundary. If the boundary is weak, the interior can store many unconstrained degrees of freedom, and the model will overfit. If the boundary is strong—for example, through attention bottlenecks, modularity, or world–model constraints—then only closure–stable structure can persist. The model is forced to extract invariants that apply beyond the training set.

The same principle operates in physics. Conservation laws, symmetries, and horizons act as boundaries that constrain what forms of structure are possible. Triadic closures that violate these boundaries cannot persist; only those compatible with them survive and propagate. Physical constants and particle identities are therefore not arbitrary but holographically determined by the boundary conditions of the universe.

Holographic boundaries thus provide the bridge between scale and meaning. Increasing capacity without tightening boundaries merely adds entropy to the interior. Increasing capacity under strong boundary constraints, however, allows more triadic closures to coexist while still converging on a small set of invariants. This is why both physical law and intelligent behavior exhibit robustness and universality: they are the fixed points of closure under constraint.

9 Implications and Predictions

If triadic closure is the primitive of organized reality, it leads to concrete implications across physics, artificial intelligence, and cognition.

Artificial intelligence. Intelligent behavior arises when boundary conditions force large interior spaces to collapse onto invariant structure. This predicts that future AI systems will improve not primarily through parameter count, but through architectures that impose strong, task–appropriate boundaries: modular networks, attention bottlenecks, world models, and explicit memory systems that create stable triadic records. Systems that maximize closure capacity under constraint should exhibit higher generalization, robustness, and sample efficiency than unconstrained models of similar size.

Quantum measurement. The triadic account of measurement implies that interference is lost when, and only when, a persistent record is formed. Quantum eraser experiments should be describable entirely in terms of the creation and destruction of triadic closures, without invoking retrocausality or observer dependence. This predicts that any experimental arrangement that prevents the stabilization of a record—even if entanglement has occurred—will restore interference.

Spacetime and gravity. If spacetime is the ledger of triadic records, then regions with high record density should correspond to regions of high curvature. This suggests a direct connection between entropy production, information storage, and gravitational effects. In particular, horizons and black holes should appear as extreme cases of closure, where records are trapped and causal access is bounded.

Particles and masses. Particles correspond to self–sustaining timelike triadic closures. Their masses and charges are invariants of the closure structures they realize. This predicts that particle properties should be derivable from purely combinatorial or geometric closure constraints, without free parameters, and that only a discrete set of such stable closures should exist.

Cognition. Thought and perception consist of the formation and manipulation of triadic closures that bind sensations, concepts, and memories into stable relations. Disorders of cognition should therefore be describable as failures of closure: either excessive frustration, in which relations cannot stabilize, or excessive rigidity, in which closures cannot be revised.

Together, these implications indicate that triadic closure is not merely a descriptive tool but a predictive framework. By identifying where and how closures form, it becomes possible to anticipate the behavior of complex systems across scales.

10 From Boundary Tension to Metric Dilation

In the triadic–closure framework, every stable fact corresponds to the formation of a persistent record. Such records arise whenever a relational configuration becomes self–maintaining, and their accumulation constitutes the informational ledger of the universe. At the largest scale, the union of all such records defines a global causal boundary, or horizon, which we denote by $\partial\Omega(t)$ with area $A(t)$. This boundary is not a material surface but the locus at which irreversible relational information is stored.

Let $S(t)$ denote the total amount of recorded information in the universe. Each new triadic closure increases S, so that

$$\frac{dS}{dt} = \sigma(t) \geq 0,$$

where $\sigma(t)$ is the global rate of record creation. By the holographic principle, the total information that can be stored within a region is bounded by the area of its boundary,

$$S(t) \leq \frac{A(t)}{4\ell_P^2}.$$

In a closure–driven universe, this bound is expected to be approached, as the system packs as much relational structure as possible into its available boundary. We therefore take

$$S(t) \approx \frac{A(t)}{4\ell_P^2}.$$

Differentiating yields

$$\frac{dA}{dt} = 4\ell_P^2 \frac{dS}{dt} = 4\ell_P^2 \sigma(t),$$

which expresses a simple but powerful result: boundary area must grow whenever irreversible records accumulate.

This growth has a direct geometric interpretation. For a spherical horizon, $A(t) = 4\pi R(t)^2$, where $R(t)$ is the effective horizon radius. The above relation implies

$$\frac{dR}{dt} = \frac{\ell_P^2}{2\pi R}\,\sigma(t).$$

Defining a cosmological scale factor $a(t) \propto R(t)$, one obtains

$$\frac{\dot{a}}{a} = \frac{\ell_P^2}{2\pi R^2}\,\sigma(t),$$

so that the expansion rate is controlled by the rate of record creation per unit boundary area. When $\sigma(t)$ is approximately constant, the expansion is exponential, reproducing a dark–energy–like regime without invoking vacuum energy.

The energetic origin of this growth lies in boundary tension. As shown in the electromagnetic case, gradients of a boundary phase field Θ store energy,

$$E_{\partial\Omega} = \int_{\partial\Omega} (\nabla\Theta)^2 \, dA.$$

Each new record increases the complexity of Θ, raising the boundary tension. To maintain holographic saturation, the boundary must then increase its area, which reduces the tension per unit area. Metric dilation is therefore the geometric response to informational stress.

Spacetime expansion in this view does not require an external arena or a repulsive force. It is the self–referential growth of the universe's relational ledger. As oscillatory dynamics generate ever more triadic closures, the holographic boundary that stores them must grow, and the induced geometry we call spacetime expands accordingly.

11 Conclusion

We have argued that triadic closure is the minimal primitive from which stable structure, meaning, and physical reality arise. Dyadic relations can express alternatives, but they cannot by themselves generate facts, identities, or invariants. Only when a relation is embedded in a triad that includes a persistent constraint or record does it become part of an ordered island capable of surviving interaction with an entropic environment.

By framing quantum measurement, spacetime geometry, particle identity, and artificial intelligence within the same closure–based logic, we obtain a unified account of why certain structures endure. Records create facts, facts create geometry, and geometry creates the stage on which further closures can occur. What we ordinarily call laws of nature, concepts, or learned representations are simply the invariants of this process: the patterns that remain unchanged under all admissible triadic projections.

This perspective dissolves several long–standing puzzles. The measurement problem is recast as the question of when a closure becomes a persistent record. The emergence of spacetime is understood as the growth of a relational ledger. The scaling behavior of modern AI systems is explained by their increasing capacity to host simultaneous closures under constraint. In each case, no appeal to observers, minds, or extraneous postulates is required.

Triadic closure therefore provides a common language for physics, cognition, and computation. It suggests that reality is not built from objects or bits, but from loops of constrained alternatives that stabilize themselves into form. Understanding and engineering such loops may be the key not only to explaining the universe we inhabit, but to constructing new forms of intelligence within it.

Foundational Position

This work does not introduce new particles, forces, or fields. It makes a more basic claim: that a minimal structural requirement for anything to exist as a stable, referable fact has been implicit in physics, computation, and cognition all along but has never been explicitly formalized.

That requirement is *triadic closure*.

A dyadic relation (A–B) can describe correlation, mapping, or superposition, but it cannot by itself produce persistence. Without a third element that encodes, constrains, or records the relation, any such pairing is vulnerable to noise, erasure, or ambiguity. Only when a relation is embedded in a closure

$$(A, B, R_{AB})$$

— where R_{AB} is a persistent relational constraint — does it become a fact.

This structure is not an interpretation layered onto physics; it is already present in its practice. A particle is not merely a state but a self-maintaining identity stabilized by fields and conservation laws. A measurement is not system plus apparatus but that pair together with an irreversible record. A bit is not 0 or 1 but a distinction held stable by a physical memory. A logical statement is not a mapping but a mapping preserved by a proof, model, or database.

Across domains, stability requires closure.

The thesis of this paper is that this triadic structure is the minimal unit of organized reality. It is the atom of facts, the seed of spacetime, and the primitive of meaning. From it follow quantum collapse, classical geometry, particle identity, and the operation of modern machine intelligence.

Seen in this light, the framework presented here is not speculative in the sense of inventing new entities. It is clarificatory: it makes explicit a structural necessity that was already embedded in the way physical and informational systems must function to exist at all.

Whenever something persists, somewhere a triadic closure is holding it in place.

Conversation as Boundary Refinement: A VectorLogic Model of Interactive Inference

Abstract

We propose that human and artificial dialogue implements a boundary-driven inference process that can be formalized as VectorLogic: constraint propagation over a quotient world. Each conversational turn refines a boundary object that determines which distinctions are meaningful relative to a question and a resolution scale. Understanding corresponds to holographic boundary completeness, not to belief states or symbol grounding. This framework explains why large language models succeed, why dialogue converges when productive, and why misunderstandings persist when boundaries remain under-specified.

---

1 Introduction

Conversation is normally treated as an exchange of beliefs, messages, or semantic content. However, in practice, dialogue functions as a constraint-tightening process: speakers negotiate definitions, assumptions, and admissible inferences until the interior space of possible interpretations collapses onto a small set of invariants.

We propose that dialogue implements a weak form of *VectorLogic* — a boundary-driven inference algorithm that operates on the quotient world induced by finite resolution and observational equivalence.

This paper formalizes that process.

---

2 Quotient Worlds and Observational Resolution

Let W be the set of microstates of a domain.
Let $\sim$ be an equivalence relation induced by the observer's resolution and interventions.

The effective world is

$$\Omega = W/{\sim}.$$

All inference occurs on Ω, not on W. Two states that cannot be distinguished or operationally separated are the same for inference.

This applies equally to physics, cognition, and dialogue.

---

3 Boundary Objects

A boundary B is a set of constraints on Ω sufficient to define what is relevant to a question Q.

Examples:

- "Classical GR only"
- "No dark energy"
- "Electron mass fixed"
- "Use everyday language"

Boundaries do not specify outcomes — they specify what is admissible.

---

4 VectorLogic

VectorLogic is the algorithm:

$$(Q, {\sim}) \xrightarrow{\text{Grip}} B \xrightarrow{\text{Boundary}} \text{Int}(B) \xrightarrow{\text{Interior}} \text{Invariants}.$$

Where:

- $Q =$ question
- $\sim =$ resolution class
- $B =$ boundary
- $\text{Int}(B) = \{x \in \Omega \mid x \models B\}$

The invariants of $\text{Int}(B)$ are what the system can conclude.

---

5 Conversation as VectorLogic

A dialogue is a sequence of triples $(Q_t, \sim_t, B_t)$.

Each turn performs:

1. User proposes $Q_t, \sim_t, C_t$
2. Agent computes $\mathrm{Int}(B_t)$
3. Agent outputs invariants $\to$ boundary refinement $B_{t+1} = B_t \cup \Delta B$

Thus:

$$B \to B' \to B'' \to \dots$$

Conversation is boundary convergence.

---

6 Holographic Adequacy

Define

$$\eta(B; Q, \sim) = 1 - \frac{L(O_{Q,\sim} \mid B)}{L(O_{Q,\sim})}.$$

This measures how much of the interior uncertainty has been eliminated by the boundary.

Understanding corresponds to $\eta \to 1$.

---

7 Weak vs Strong VectorLogic

Property	**Weak (LLMs)**	**Strong (IC agents)**
Boundary	implicit	explicit
Closure	heuristic	formal operators
Memory	context window	invariant ledger
Stopping	satisfaction	η-saturation

LLMs approximate VectorLogic but do not track closure explicitly.

---

8 Why Dialogue Works

Productive conversation converges because:

- ambiguity triggers repair
- contradictions force refinement
- examples prune interiors
- metaphors test boundary consistency

All raise η.

Misunderstanding is low η.

---

9 Implications for AI

To build genuinely reasoning systems:

- boundaries must be explicit objects
- invariants must be stored
- adequacy must be measurable
- resolution must be negotiable

This yields interpretable, steerable cognition.

---

10 Conclusion

Conversation is not symbol exchange.
It is holographic boundary refinement under finite bandwidth.

VectorLogic explains why humans, science, and LLMs converge when they do — and why they fail when they do not.

Understanding is not in the words.
It is in the boundary.

Logic as Holographic Inference:
Boundary Constraints and the Structure of Understanding

Mikael Bergkvist [Coauthor / Collaborator]

Abstract

We present a reformulation of logic as boundary-driven inference rather than rule-based reasoning. In this framework, understanding a problem consists of three operations: (i) *gripping* the problem by identifying its domain and observational resolution, (ii) specifying a set of immutable constraints (the *boundary*), and (iii) deriving the admissible configurations satisfying those constraints (the *interior*). We formalize an *epistemic holography* principle: when the boundary is complete relative to a question and resolution, the interior adds no independent explanatory content; it is an inference-unpacking of boundary information. This perspective suggests that many persistent "mysteries" arise when explanatory frameworks privilege time-directed propagation before making constraints explicit. As a flagship case study, we show how a constraint-first formulation of the quantum harmonic oscillator organizes spectra, nodal structure, and measurement updates as consequences of a single boundary specification, with probabilistic weights emerging from Hilbert-space projection geometry under explicit assumptions. We then outline (as hypotheses) boundary-first reinterpretations of cosmological acceleration and cognitive binding, and derive design implications for AI: internal communication bottlenecks and constraint-native learning can force abstraction as an invariant under lossy translation.

Status Legend.

- **D (Derived):** follows formally from stated assumptions.
- **R (Reformulation):** empirically equivalent, primarily conceptual/structural.
- **H (Hypothesis):** proposed extension; testable but not established here.

Contents

1 Introduction

1.1 The problem with evolution-first explanation

Many explanations in physics, mathematics, and cognitive science are presented as time-directed narratives:

$$\text{“this causes that”} \rightarrow \text{“then this happens”} \rightarrow \text{“therefore we observe } X\text{”}.$$

Such narratives are computationally effective, but they can obscure a more fundamental question: *which constraints make the phenomenon necessary?* In several prominent domains, evolution-first formalisms appear to generate persistent interpretational residue—“mysteries” that resist resolution for decades.

1.2 Motivating examples

We briefly list canonical examples (expanded later):

1. **Quantum measurement:** why do observations correspond to definite outcomes?
2. **Wigner's puzzle:** why is mathematics so effective at describing the physical world?
3. **Consciousness / binding:** how do distributed processes yield unified experience?
4. **Cosmology / acceleration:** why does expansion appear to accelerate (“dark energy”)?

1.3 Thesis: logic as boundary-driven inference

We propose that logic, in its most general (cross-domain) form, is not rule-following but *boundary inference*. Understanding a problem consists of three operations:

1. **Grip:** specify the domain, resolution, and what is being asked.
2. **Boundary:** identify immutable constraints that define the problem's identity.
3. **Interior:** derive the admissible configurations forced by those constraints.

A central claim is that when a boundary specification is *complete* relative to a question and observational resolution, the interior is an *epistemically holographic* unpacking of boundary information.

1.4 Roadmap

Section 2 develops the formal framework, including an Axiom of Finite Correspondence and a diagnostic for “false mysteries.” Section 3 applies the framework to quantum mechanics with a harmonic-oscillator case study. Section 4 outlines extensions to cosmology and consciousness (as hypotheses) and clarifies what is claimed vs. conjectured. Section 5 discusses AI implications and proposes constraint-native architectural principles. Section 6 addresses objections. Section 7 concludes.

2 Formal Framework

2.1 Basic objects: world, observer, and resolution

Definition 1 (World state space)**.** *Let W be a set of microstates (“world states”) in some regime of interest.*

Definition 2 (Actions and interventions)**.** *Let $\mathcal{A}$ be a set of admissible interventions. Each $a \in \mathcal{A}$ induces a world update map*

$$T_a : W \to W.$$

Definition 3 (Outcome map and observational resolution)**.** *Let O be a set of detector outcomes and let $m : W \to O$ be a measurement/readout map. Define an observational equivalence relation $\sim$ on W by*

$$w_1 \sim w_2 \iff \forall a \in \mathcal{A} : \; m(T_a(w_1)) = m(T_a(w_2)).$$

The induced observable quotient is $\Omega := W/\sim$.

2.2 Boundary and interior

Definition 4 (Boundary constraint)**.** *A* boundary *is a set B of constraints that cannot be violated without (i) contradiction, (ii) disagreement with empirical regularities assumed in-regime, or (iii) changing the identity of the question under study.*

Definition 5 (Interior)**.** *Given a constraint space C (domain of candidate configurations) and boundary set $B \subseteq C$, define the* interior *(feasible set) as*

$$\mathrm{Int}(B) := \{x \in C : \; x \textit{ satisfies all } b \in B\}.$$

Definition 6 (Completeness relative to a question)**.** *Fix a question Q and resolution $\sim$. A boundary set B is* complete (relative to $Q, \sim$) *if (up to gauge equivalence)* $\mathrm{Int}(B)$ *determines all observables relevant to Q at resolution $\sim$ and admits no additional independent degrees of freedom that change those observables.*

2.3 Epistemic holography

Definition 7 (Epistemic holography)**.** *An inference problem $(C, B, Q, \sim)$ is* epistemically holographic *if, relative to Q and $\sim$, any "bulk" augmentation of the interior adds no predictive power beyond the boundary specification. Informally:* the interior is boundary-unpacking.

Remark 1. *This notion is epistemic/structural, not a direct claim about physical area laws. Physical holography (e.g. AdS/CFT, black-hole entropy bounds) may be viewed as one instantiation, but is not required for the logical framework.*

2.4 Axiom of Finite Correspondence and Wigner's puzzle

Axiom 1 (Finite Correspondence)**.** *In any observational regime with finite resources, there exists an observational equivalence $\sim$ such that the induced observable quotient $\Omega = W/\sim$ is finite:*

$$|\Omega| < \infty.$$

Definition 8 (Induced observable dynamics)**.** *Each $T_a : W \to W$ induces a well-defined map $\bar{T}_a : \Omega \to \Omega$ via $\bar{T}_a([w]) := [T_a(w)]$. Let M be the transformation monoid generated by $\{\bar{T}_a\}_{a \in \mathcal{A}}$ under composition.*

Definition 9 (Observer representation)**.** *Let an observer have an internal state set Σ and internal update operators $\tau_a : \Sigma \to \Sigma$. A representation is a coding map $\varphi : \Omega \to \Sigma$ such that for all $a \in \mathcal{A}$,*

$$\varphi \circ \bar{T}_a = \tau_a \circ \varphi.$$

Proposition 1 (Homomorphic inevitability (D)). *Under Axiom 1, any observer that consistently predicts the effect of interventions on distinguishable outcomes implements a monoid homomorphism*

$$\Phi : \mathsf{M} \to \Sigma^{\Sigma} \quad \textit{such that} \quad \Phi(\bar{T}_a) = \tau_a \textit{ and } \Phi(f \circ g) = \Phi(f) \circ \Phi(g).$$

Remark 2. *This reframes Wigner's "unreasonable effectiveness" as a structural necessity: finite distinguishability plus compositional intervention structure forces a structure-preserving internal model.*

Proposition 2 (Isomorphism onto image (D)). *If* $\mathrm{Ker}(\Phi) = \{\mathrm{id}\}$ *then* $\mathsf{M} \cong \mathrm{Im}(\Phi)$.

2.5 The three operations: Grip, Boundary, Interior

Definition 10 (Operation 1: Grip). Gripping *a problem means specifying (i) the domain/constraint space C, (ii) the observational resolution $\sim$ (what counts as the same), and (iii) the question Q (which observables matter).*

Definition 11 (Operation 2: Boundary). Boundary definition *means selecting a set B of immutable constraints sufficient for completeness relative to $(Q, \sim)$.*

Definition 12 (Operation 3: Interior). Interior derivation *means computing or characterizing* $\mathrm{Int}(B)$ *and its invariants, by optimization, deduction, or algorithmic constraint satisfaction.*

2.6 Information inequality (epistemic)

Proposition 3 (Boundary compression bound (R)). *If B is complete relative to $(Q, \sim)$, then any minimal description of* $\mathrm{Int}(B)$ *sufficient to answer Q at resolution $\sim$ has complexity bounded (up to constants) by the complexity of B.*

2.7 Diagnostic: false mysteries vs genuine mysteries

Definition 13 (Mystery diagnostic). *Given a "mystery" M in an evolution-first formulation, apply:*

1. *Identify Q and $\sim$ (grip).*
2. *Make constraints explicit and attempt a boundary-value / constraint-satisfaction formulation.*
3. *If M vanishes when constraints are complete, classify as* false mystery *(formulation artifact).*
4. *If M persists due to missing constraints or missing empirical boundaries, classify as* genuine mystery*.*

2.8 Boundary Completeness and Self-Validation

A key feature of boundary-driven logic is that it embeds a self-testing mechanism. A boundary specification B is not assumed correct a priori; it is validated or falsified by the interior it generates.

Definition 14 (Completeness Criterion). *A boundary B is complete relative to a question Q and observational resolution $\sim$ if and only if the derived interior* $\mathrm{Int}(B)$ *accounts for all observables relevant to Q at resolution $\sim$, and admits no additional degrees of freedom capable of altering those observables.*

Operationally, completeness is assessed by computing $\mathrm{Int}(B)$ and testing whether any mismatch remains between its predictions and empirical or logical constraints defining Q. Any residual discrepancy signals that B is incomplete or improperly specified, necessitating boundary refinement.

This induces an iterative refinement loop:

$$B \;\to\; \mathrm{Int}(B) \;\to\; \mathrm{Test}(Q, \sim) \;\to\; B',$$

which converges when interior variation becomes purely gauge with respect to Q. At convergence, the interior is epistemically holographic: no bulk augmentation adds explanatory power.

This mechanism underlies the "mystery diagnostic" proposed in this work: if a purported mystery vanishes upon boundary completion, it is a formulation artifact rather than a fundamental indeterminacy.

3 Application: Quantum Mechanics

3.1 Standard evolution-first formulation (context)

The Schröodinger equation

$$i\hbar \frac{\partial \psi}{\partial t} = \hat{H}\psi$$

provides a propagator: given an initial condition, it yields a time-ordered state update. This is computationally powerful, but interpretational issues often appear when evolution is taken as ontologically primary.

3.2 Constraint-first formulation: eigenstate hierarchy

We consider a time-symmetric constraint formulation for stationary structure. Let $\psi \in \mathcal{H}$ (a Hilbert space) with normalization $\langle \psi | \psi \rangle = 1$.

Definition 15 (Tension functional)**.** *Define the functional*

$$\mathcal{T}[\psi] := \langle \psi | \hat{H}^2 | \psi \rangle,$$

subject to normalization and orthogonality constraints against previously found states.

Definition 16 (Nested minimization (eigenstate hierarchy))**.** *Define ψ_0 as a minimizer of $\mathcal{T}[\psi]$ under $\langle \psi | \psi \rangle = 1$. Then define ψ_n as a minimizer under $\langle \psi | \psi \rangle = 1$ and $\langle \psi | \psi_k \rangle = 0$ for all $k < n$.*

Proposition 4 (Spectrum recovery (D))**.** *Under standard spectral assumptions on $\hat{H}$ (self-adjointness, discrete spectrum in the bound subspace), the nested minimization yields the eigenbasis ordering and*

$$\mathcal{T}[\psi_n] = E_n^2 \quad \textit{with} \quad \hat{H}\psi_n = E_n \psi_n$$

(up to sign conventions for E_n).

3.3 Harmonic oscillator case study (sketch)

For the 1D harmonic oscillator, one may use trial families to show that (i) optimal width emerges from balancing kinetic vs potential terms, and (ii) orthogonality constraints force nodal structure (Hermite polynomials).

Ground state trial. Let

$$\psi_\sigma(x) = \left(\frac{1}{\pi\sigma^2}\right)^{1/4} \exp\left(-\frac{x^2}{2\sigma^2}\right).$$

Compute $\mathcal{T}(\sigma)$ and minimize over σ to obtain $\sigma^2 = \hbar/(m\omega)$ and $\mathcal{T}^{(0)}_{\min} = E_0^2$.

Excited states. Impose parity and orthogonality constraints via trial families of the form

$$\psi_n(x) \propto p_n(x) \exp\left(-\frac{x^2}{2\sigma^2}\right),$$

with p_n polynomial and orthogonality to all lower states. Optimization yields Hermite structure as an emergent optimal packing under exclusion constraints.

3.4 Measurement update as constraint addition

In evolution-first accounts, "collapse" is an additional postulate. In the boundary-first view, measurement corresponds to adding a new constraint to B and re-deriving the interior.

Definition 17 (Constraint addition)**.** *Let B be a boundary set defining an admissible set* $\mathrm{Int}(B)$*. A measurement outcome is modeled as an added constraint $b^\star$; the updated interior is*

$$\mathrm{Int}(B') \quad \textit{where} \quad B' := B \cup \{b^\star\}.$$

3.5 Holographic Adequacy as a Metric of Understanding

A central advantage of the boundary–interior framework is that it admits a *quantitative measure of understanding.* Rather than treating understanding as subjective or binary, we characterize it as the degree to which a boundary specification accounts for all relevant observables without introducing independent interior degrees of freedom. This motivates a holographic adequacy metric.

3.5.1 Effective Observables and Resolution

Let Q denote a question posed about a system, and let $\sim$ denote the observational resolution at which Q is evaluated. Let $\mathcal{O}_{Q,\sim}$ be the set of observables relevant to Q at resolution $\sim$.

Two interior configurations are considered equivalent if they induce the same values on $\mathcal{O}_{Q,\sim}$. We therefore define all interior quantities *modulo* this equivalence relation.

3.5.2 Boundary and Interior Description Length

Let B be a boundary specification (a finite set of immutable constraints), and let $\mathrm{Int}(B)$ denote the interior inferred from B.

We define:

- $L(B)$: the description length of the boundary constraints,
- $L(\mathcal{O}_{Q,\sim})$: the description length of the observables,
- $L(\mathcal{O}_{Q,\sim} \mid B)$: the additional description length required to reproduce $\mathcal{O}_{Q,\sim}$ given boundary B.

Description length may be instantiated using minimum description length (MDL), Bayesian evidence, or any fixed encoding scheme appropriate to the domain. No appeal to uncomputable Kolmogorov complexity is required.

3.5.3 Holographic Adequacy

Definition 18 (Holographic Adequacy)**.** *The holographic adequacy of boundary B relative to question Q and resolution $\sim$ is defined as*

$$\eta_{\text{holo}}(B; Q, \sim) \;=\; 1 - \frac{L(\mathcal{O}_{Q,\sim} \mid B)}{L(\mathcal{O}_{Q,\sim})}.$$

Interpretation:

- $\eta_{\text{holo}} \to 1$: boundary nearly determines all relevant observables; understanding is near-complete.
- $\eta_{\text{holo}} \ll 1$: substantial explanatory bulk remains; boundary is incomplete.

This metric directly captures the explanatory sufficiency of the boundary without conflating representational verbosity with physical content.

3.5.4 Residual Interior Freedom

To distinguish incomplete boundaries from mere representational redundancy, we define the effective interior degrees of freedom:

$$\text{DOF}_{Q,\sim}(\text{Int}(B)) \;=\; \dim(\text{Int}(B)/\!\sim),$$

i.e., the number of interior variations that change observables relevant to Q at resolution $\sim$.

A refined adequacy score incorporating residual interior freedom is then:

$$\eta'_{\text{holo}} \;=\; 1 - \frac{L(\mathcal{O}_{Q,\sim} \mid B) + \lambda\, \text{DOF}_{Q,\sim}(\text{Int}(B))}{L(\mathcal{O}_{Q,\sim})},$$

where λ sets the penalty weight for unexplained adjustable structure.

3.5.5 Diagnostic Interpretation

- If $L(\mathcal{O}_{Q,\sim} \mid B)$ is large, the boundary is incomplete: additional constraints are required.
- If $\text{DOF}_{Q,\sim}$ is large, the formulation contains hidden bulk degrees of freedom not fixed by the boundary.
- If both are small, the formulation is holographically adequate.

This provides an operational mystery diagnostic: apparent paradoxes correspond to low holographic adequacy and should diminish as η_{holo} increases under boundary refinement.

3.5.6 Iterative Boundary Refinement

Understanding becomes an optimization process:

$$B_0 \;\to\; \text{Int}(B_0) \;\to\; \eta_0 \;\to\; B_1 \;\to\; \eta_1 \;\to\; \cdots$$

Boundary modifications are retained only if they increase holographic adequacy. Convergence occurs when η_{holo} saturates or no further boundary compression is possible at the chosen resolution.

3.5.7 Example: Quantum Harmonic Oscillator

In the standard Schrödinger formulation, the boundary specification (evolution equation plus normalization) leaves substantial interior structure unaccounted for, yielding low holographic adequacy. In contrast, the constraint-based formulation presented in this work derives eigenvalues, eigenstates, uncertainty relations, and measurement behavior directly from a minimal boundary, achieving significantly higher η_{holo}.

3.5.8 Implications

This metric:

1. provides an objective measure of explanatory completeness,
2. allows principled comparison between formulations,
3. distinguishes genuine open problems from formulation artifacts,
4. supplies a quantitative test for AI generalization and abstraction.

In this sense, understanding acquires a metric structure: proximity to holographic adequacy.

Remark 3 (Born weights: status)**.** *Squared projection norms arise naturally from Hilbert-space geometry. Interpreting these as probabilities typically requires an additional interpretational axiom set (e.g. Gleason-type assumptions, symmetry/typicality, or decision-theoretic postulates). We treat this as* ***R*** *in the present manuscript unless explicitly assumed.*

3.6 Summary: what changes and what does not

- **Predictions:** unchanged (same spectral structure, same measurement statistics under standard assumptions).
- **Explanatory load:** shifted from time-directed narration to constraint completeness.
- **Mystery diagnostic:** "collapse" becomes a boundary update; remaining questions reduce to boundary completeness and interpretational assumptions.

4 Extended Applications

4.1 Cosmology: boundary-first reinterpretation of acceleration (H)

Standard cosmology often treats acceleration via effective bulk terms (e.g. Λ). A boundary-first hypothesis is that some features attributed to bulk substances can be re-expressed as consequences of boundary information growth under finiteness constraints.

Hypothesis. In a finite correspondence regime, the relevant macro-observables are functions of boundary-encoded degrees of freedom; apparent acceleration may be modeled as an emergent parameter describing boundary growth as seen from an interior reconstruction.

Empirical handles. Any such hypothesis must specify (i) an effective equation of state or equivalent observational signature, and (ii) at least one distinguishable prediction relative to ΛCDM in precision surveys. We outline candidate signatures without claiming a completed derivation here.

4.2 Consciousness: binding as constraint consistency (H)

The binding problem can be reframed as: how does a system select a unified interior when multiple distributed constraint sources must be jointly satisfied?

Hypothesis. Conscious experience corresponds to the internally reconstructed interior $\mathrm{Int}(B_t)$ of a time-local active boundary B_t (the "active constraint network"). Unity is a consistency property: incompatible constraints produce rivalry/fragmentation; compatible constraints yield a single integrated reconstruction.

Scope. This addresses *binding/unity* structurally. Mapping basis choice to felt qualitative character (qualia) remains an open problem, but the framework eliminates the need for an additional "binding mechanism" beyond constraint compatibility.

4.3 Wigner's puzzle revisited (D/R)

Axiom 1 plus Proposition 1 implies that mathematical structure is forced in any stable observer: finite distinguishability and compositional interventions induce algebra, and predictive modeling requires homomorphisms preserving that algebra.

5 Implications for AI

5.1 Evolution-first learning and its limits

Gradient-based deep learning can be viewed as an evolution-first optimization process (trajectory-first) that may optimize away internal communication constraints. In this view, abstraction is optional: a byproduct rather than a necessity.

5.2 Constraint-native learning: design principle (H)

We propose that rapid generalization can be promoted by architectures that are *constraint-native*:

- multiple specialized "islands" with incompatible representations,
- explicit communication bottlenecks between islands,
- a global objective that requires coordination across islands,
- an update rule that preserves invariants under lossy translation.

5.3 IC as abstraction-forcing translation

Let islands i have internal state spaces Σ_i and translation channels $T_{i\to j}$ with limited capacity. Define an abstraction as an invariant surviving translation across the island family:

$$\mathcal{A} := \bigcap_{i,j} \mathrm{Inv}(T_{i\to j}).$$

This operationalizes "abstraction = invariance under lossy translation."

5.4 Testable predictions (H)

Compared to monolithic models with similar parameter counts, constraint-native systems should show:

1. sharper transfer learning (fewer samples for new tasks),
2. improved robustness under distribution shift (violation detection),
3. more interpretable intermediate representations (explicit invariants),
4. phase-transition-like capability jumps when compatibility thresholds are crossed.

6 Objections and Responses

6.1 "This is just reformulation, not new physics"

Response: often correct and intended. The principal claim is about explanatory structure: reformulations can eliminate auxiliary postulates by making constraints explicit. Conceptual progress may occur even when empirical content is preserved.

6.2 "Born rule does not drop out"

Response: agreed in general. Projection geometry is immediate; probability interpretation requires explicit assumptions. We mark this boundary clearly and treat it as a separate axiom layer when needed.

6.3 "Holography is being misused"

Response: we distinguish *epistemic holography* (definition in Section 2) from physical holography (area laws, AdS/CFT). The former is a property of inference problems; the latter is a claim about physical degrees of freedom.

6.4 "Cosmology and consciousness are speculative"

Response: yes; they are labeled **H**. They are included to show how the framework generates testable research programs, not as completed derivations.

7 Conclusion

We argued that logic is most fruitfully understood as boundary-driven inference: grip the problem, specify immutable constraints, and derive interiors that satisfy them. When the boundary is complete relative to a question and observational resolution, the interior is epistemically holographic: it adds no independent explanatory content. This re-weights explanation away from time-directed narratives toward constraint completeness, and reframes Wigner's puzzle as homomorphic inevitability under finite correspondence. Quantum mechanics provides a flagship demonstration that constraint-first formalisms can concentrate explanatory power and clarify measurement updates as boundary additions. Extensions to cosmology and consciousness remain hypotheses but illustrate a unifying research program. Finally, we suggest that AI generalization can be accelerated by constraint-native architectures that force abstraction as invariance under lossy translation.

Future work. (i) a systematic catalog of false mysteries, (ii) full technical derivations and assumption sets for probability interpretation, (iii) concrete AI benchmarks for constraint-native architectures, and (iv) observationally distinguishable cosmological signatures of boundary-first acceleration models.

A Appendix A: Proof details

A.1 Proof of Proposition 1 (Homomorphic Inevitability)

Statement (Proposition 1). Let W be the set of *world-states* and let $\sim$ be an observational equivalence relation induced by a fixed observational resolution (or coarse graining) for an observer G_A. Let $\pi : W \to \Omega$ be the quotient map onto

$$\Omega \;:=\; W/\!\sim \;=\; \{[w]_\sim : w \in W\}, \qquad \pi(w) := [w]_\sim.$$

Assume the observer's representation $\Phi : W \to \mathcal{R}$ is $\sim$*-invariant* (i.e. respects indistinguishability):

$$\forall w, w' \in W : \quad w \sim w' \implies \Phi(w) = \Phi(w'). \tag{1}$$

Then Φ factors uniquely through the quotient:

$$\exists!\, \bar{\Phi} : \Omega \to \mathcal{R} \quad \text{s.t.} \quad \Phi \;=\; \bar{\Phi} \circ \pi.$$

Moreover, if the system admits a family of admissible interventions whose composition induces an algebraic structure on Ω (defined below), and if Φ is required to respect empirical composition (closure under intervention), then $\bar{\Phi}$ must be a homomorphism with respect to that structure. This yields the *homomorphic inevitability*: the observer's mathematical structure is forced by the quotient geometry Ω.

Interpretive note (Wigner loop closure). The key mechanism is the quotient $\Omega = W/\sim$: once an observer can only distinguish equivalence classes, any stable predictive representation must *descend* to Ω. If the observer also tracks how experimental actions compose, that compositional structure must also descend, forcing Φ to respect it—i.e. become a homomorphism. Mathematics is then not an optional overlay but the inevitable language of the quotient.

A.1.1 A.1.1 Construction: the quotient and the projection map

Definition 19 (Observational equivalence)**.** *Fix an observer G_A with resolution $\sim$. Define $w \sim w'$ iff G_A cannot distinguish w and w' by any admissible observation at that resolution.*

$\sim$ is assumed to be an equivalence relation (reflexive, symmetric, transitive). The quotient set $\Omega = W/\sim$ is the set of equivalence classes $[w]_\sim$. The canonical projection is $\pi : W \to \Omega$ given by $\pi(w) = [w]_\sim$.

A.1.2 A.1.2 Lemma 1: factorization through the quotient

[Quotient factorization / universal property] Let $\Phi : W \to \mathcal{R}$ satisfy $\sim$-invariance (1). Then there exists a unique $\bar{\Phi} : \Omega \to \mathcal{R}$ such that $\Phi = \bar{\Phi} \circ \pi$.

Proof. Define $\bar{\Phi} : \Omega \to \mathcal{R}$ by

$$\bar{\Phi}([w]_\sim) \ := \ \Phi(w).$$

We must show this is well-defined. If $[w]_\sim = [w']_\sim$, then $w \sim w'$, hence $\Phi(w) = \Phi(w')$ by (1). Therefore $\bar{\Phi}$ does not depend on the choice of representative.

Now for any $w \in W$,

$$(\bar{\Phi} \circ \pi)(w) = \bar{\Phi}([w]_\sim) = \Phi(w),$$

so $\Phi = \bar{\Phi} \circ \pi$.

Uniqueness: if $\Phi = \tilde{\Phi} \circ \pi$ for some $\tilde{\Phi} : \Omega \to \mathcal{R}$, then for all $[w]_\sim \in \Omega$,

$$\tilde{\Phi}([w]_\sim) = \tilde{\Phi}(\pi(w)) = \Phi(w) = \bar{\Phi}([w]_\sim),$$

so $\tilde{\Phi} = \bar{\Phi}$. □

Consequence. Lemma A.1.2 already closes the first half of Wigner's loop: *any* representation stable under observational equivalence cannot be a free function on W; it is forced to be a function on Ω.

A.1.3 A.1.3 Adding dynamics/experiments: induced structure on Ω

To get *homomorphism* rather than mere factorization, we encode the empirical fact that the observer does not only see states, but also performs *interventions* and observes how they compose.

Definition 20 (Admissible interventions)**.** *Let $\mathcal{U}$ be a set of admissible interventions. Each $u \in \mathcal{U}$ acts as a (possibly partial) function $u : W \to W$. Composition is function composition: $(u \circ v)(w) = u(v(w))$ whenever defined.*

The observer cannot access $w \in W$ directly, only $\pi(w) \in \Omega$. Thus we need interventions to be *compatible* with the observational equivalence.

Definition 21 ($\sim$ as a congruence for interventions)**.** *We say $\sim$ is a congruence with respect to $\mathcal{U}$ if*

$$\forall u \in \mathcal{U}, \ \forall w, w' \in W : \quad w \sim w' \ \Rightarrow \ u(w) \sim u(w'). \tag{2}$$

[Induced action on the quotient] If (2) holds, then each $u \in \mathcal{U}$ induces a well-defined map $\hat{u} : \Omega \to \Omega$ by

$$\hat{u}([w]_\sim) \ := \ [u(w)]_\sim.$$

Moreover, $\widehat{u \circ v} = \hat{u} \circ \hat{v}$ whenever $u \circ v$ is defined.

Proof. Well-definedness: if $[w]_\sim = [w']_\sim$ then $w \sim w'$, and by (2) we have $u(w) \sim u(w')$, so $[u(w)]_\sim = [u(w')]_\sim$.

Composition: for any $[w]_\sim \in \Omega$,

$$\widehat{u \circ v}([w]_\sim) = [(u \circ v)(w)]_\sim = [u(v(w))]_\sim = \hat{u}([v(w)]_\sim) = \hat{u}(\hat{v}([w]_\sim)).$$

□

Lemma A.1.3 says: *the algebra of experiments descends to the quotient.* The observer's empirical world is therefore $(\Omega, \hat{\mathcal{U}})$.

A.1.4 A.1.4 Lemma 2: predictive composition forces homomorphism

We now formalize the "requirement on Φ". A minimal operational requirement is that Φ respects the empirical composition of interventions up to the observer's codomain structure.

Definition 22 (Operational representation)**.** *Let $\mathcal{R}$ be a set equipped with an action of interventions $\rho : \mathcal{U} \to \mathrm{End}(\mathcal{R})$, i.e. each $u \mapsto \rho(u)$ is a function $\rho(u) : \mathcal{R} \to \mathcal{R}$ such that $\rho(u \circ v) = \rho(u) \circ \rho(v)$ whenever $u \circ v$ is defined. We say $\Phi : W \to \mathcal{R}$ is an operational representation if*

$$\forall u \in \mathcal{U} : \quad \Phi(u(w)) \; = \; \rho(u)(\Phi(w)) \quad \textit{whenever } u(w) \textit{ is defined.} \tag{3}$$

[Descent intertwines induced actions] Assume $\sim$-invariance (1) and congruence (2). Let $\bar{\Phi} : \Omega \to \mathcal{R}$ be the unique factor from Lemma A.1.2. Then for all $u \in \mathcal{U}$,

$$\bar{\Phi}(\hat{u}(\omega)) \; = \; \rho(u)(\bar{\Phi}(\omega)) \qquad \forall \omega \in \Omega,$$

i.e. $\bar{\Phi}$ is a homomorphism (an intertwiner) between the induced action on Ω and the representation action on $\mathcal{R}$.

Proof. Let $\omega = [w]_\sim$. Then using Lemma A.1.3 and (3),

$$\bar{\Phi}(\hat{u}(\omega)) = \bar{\Phi}([u(w)]_\sim) = \Phi(u(w)) = \rho(u)(\Phi(w)) = \rho(u)(\bar{\Phi}([w]_\sim)) = \rho(u)(\bar{\Phi}(\omega)).$$

□

A.1.5 A.1.5 Conclusion: the homomorphic inevitability

Combining Lemma A.1.2 and Lemma A.1.4 yields:

1. (Quotient forcing) $\sim$-stable representations *must* factor through $\pi : W \to \Omega$.
2. (Structure forcing) if interventions are congruent with $\sim$, their composition descends to Ω; any predictive representation respecting empirical composition must intertwine these actions—hence must be a homomorphism relative to the induced operational structure on Ω.

Formal closure on Wigner. Wigner's "unreasonable effectiveness" becomes a structural inevitability: the observer's mathematics is not "mysteriously matched" to the world; it is the minimal homomorphic image of the quotient world $\Omega = W/\sim$ under the requirement that predictions be invariant under observational indistinguishability and consistent under composition of interventions.

Where finiteness enters. If Ω is finite (or effectively finite at resolution $\sim$), then the observer's induced operational structure is necessarily discrete and compressible. In that regime, algebraic regularities (symmetries, invariants, closure laws) are not optional modeling choices but the most efficient representations of the quotient dynamics; they are forced by the requirement that $\bar{\Phi}$ be stable, compositional, and resolution-correct.

□

References