

The Geometry of Closure

*Admissibility, Obstruction,
and Ontological Enlargement*

Flyxion

Independent Researcher

June 3, 2026

Working draft.

Preface

This book did not begin with geometric measure theory or quantum Bayesian inference. It began with a pattern.

Across a sequence of projects developed under the names RSVP, Spherepop, CLIO, TARTAN, Simulated Agency, and the Projection Crisis essays, the same structural failure kept appearing in different contexts and under different local names. A representation would prove locally adequate, generating reliable predictions and useful intuitions within a familiar regime. Then something would appear at the boundary of that regime which the representation could not absorb. The representation did not become false. It became insufficient. The map survived. Recovery failed.

The framework developed in this book is an attempt to say precisely what that means. The central concepts are closure, recovery, obstruction, and ontological commitment. They crystallized from repeated independent encounters with the same failure mode across domains that have no obvious connection to one another. That genealogy matters: a framework designed to fit selected examples is always vulnerable to the charge of engineering. A framework that emerged elsewhere and was then found to apply has different evidential status.

Two terminological notes. *Collapse* is used throughout in a technical sense: the loss of recovery while embedding

is retained. The old theory survives; it becomes a regime. *Obstruction* is used in the sense already standard in topology, category theory, and PDE: an object that certifies, by its existence, that a certain recovery is impossible.

Three philosophical themes recur across the technical chapters, and the reader may find it useful to name them at the outset.

First: continuation is the primitive concept, not objects. The central question of the framework is never what a system *is* but whether an admissible continuation of it *exists*. Collapse is defined by the disappearance of admissible continuation; enlargement is defined by its restoration.

Second: ontologies are residues, not foundations. The objects of a theory appear fundamental only because they survive the dynamics used to generate them. Space, probability distributions, semantic scopes, and mathematical objects themselves are persistence structures: what remains invariant under repeated application of governing operations.

Third: structure without blueprint. Area-minimizing currents do not know the surface they are attempting to become. Bayesian posteriors do not know the parameter they will converge toward. Scope graphs do not contain their future resolutions. Global structure emerges from local admissibility conditions, not from pre-specified endpoints. Ontological enlargement is not the discovery of a hidden plan. It is the discovery that the admissibility conditions were richer than previously recognized.

Flyxion

Independent Researcher

Contents

Preface	i
I Closure and Recovery: The Framework	1
1 The Primitive Concept	2
1.1 What This Book Is About	2
1.2 Continuation as Primitive	3
1.3 Closure Operators and Dynamic Closure . .	4
1.4 Embedding and Recovery	6
1.5 Strong and Weak Survival	6
1.6 Invariant Systems and Collapse	8
1.7 Obstruction Theory	9
1.8 Closure as Reflection: A Categorical View . .	10
1.9 The Six-Stage Pattern	11
II The Geometry Case: A Completed Col- lapse	13
2 Functionals and the First Enlargement	14
2.1 The Brachistochrone and the Birth of Functionals	14
2.2 The Euler-Lagrange Equation	15
2.3 The Plateau Problem and Minimal Surfaces .	16

2.4	Bernstein's Theorem and the Illusion of Universality	17
3	The Simons Cone and Obstruction	18
3.1	The Construction	18
3.2	Dimension as Control Parameter	19
3.3	The Cone as Obstruction	19
3.4	Smoothness Becomes a Regime	20
4	Geometric Measure Theory as Enlarged Ontology	22
4.1	The Closure Failure of Smooth Surfaces . . .	22
4.2	Currents: Closure Restored	23
4.3	Tangent Cones and Local Models	24
4.4	The GMT Case: Five Claims	25
III	The Inference Case: An Ongoing Collapse	26
5	Classical Bayesian Inference and Its Geometry	27
5.1	Inference as Dynamical Process	27
5.2	Kullback-Leibler Divergence and the Fisher Metric	28
5.3	The Data Processing Inequality and No Epistemic Harm	28
5.4	Asymptotic Frequentism	29
5.5	Schrödinger Bridges and Process Updating .	29
6	Quantum Inference and the Prospective Collapse	31
6.1	Density Matrices and Noncommutativity . .	31

6.2	The Petz Recovery Map	32
6.3	A Worked Petz Recovery Example	32
6.4	The Invariant System and What Collapses . .	34
6.5	Established versus Prospective Collapse . . .	35
IV The Formal Case: Spherepop as Abstract Model		36
7	Scope Graphs and Semantic Topology	37
7.1	The Spherepop Calculus: Definitions	37
7.2	Scope Graphs	38
7.3	Collapse Cycles as Obstructions	39
7.4	Minimal Obstructions	39
7.5	Semantic Entropy and the Obstruction Measure	40
7.6	The Collapse Quotient and Irreversibility . .	41
7.7	Spherepop as Formal Model	43
V Diagnostic Application: Foundation Models		45
8	The Token Ontology as an Invariant System	46
8.1	The Token Ontology in the Framework	46
8.2	Candidate Enlarged Ontologies	47
8.3	Candidate Obstructions	48
8.4	Accessibility and Phenomenology	49
8.5	Framework Predictions	49

VI	General Theory and Conclusions	51
9	A General Theory of Ontological Enlargement	52
9.1	The Central Results	52
9.2	Comparative Table	53
9.3	Open Problems	55
10	Conclusions	56
10.1	The Central Claim, Evaluated	56
10.2	Structure Without Blueprint	57
10.3	Obstructions Do Not Create Enlargement . .	57
	References	60

Part I

Closure and Recovery: The Framework

Chapter 1

The Primitive Concept

A map may be accurate in every detail it shows and still fail to show the detail that matters most.

1.1 What This Book Is About

Collapse preserves ontological embedding but not ontological recovery.

That sentence is the thesis of this book. Everything that follows is an attempt to make it precise, to test it against the historical record, and to ask what it implies about mathematical knowledge more broadly.

The sentence requires three concepts. *Ontological embedding*: the old objects survive as a recognizable subclass of a larger framework. *Ontological recovery*: the governing dynamics of the larger framework remain confined to that subclass. *Collapse*: the event by which the invariant system of the old framework becomes globally unsatisfiable, forcing enlarge-

ment.

None of these is standard terminology. They are introduced as names for a structural phenomenon that keeps appearing across mathematics under different local descriptions. The phenomenon is this. A theory develops an ontology, a preferred class of objects and a set of governing operations, and that ontology works reliably within a certain regime. Then the governing operations, applied without restriction, produce objects the ontology cannot contain. The ontology must be enlarged. But enlargement is not symmetric. The old objects embed into the new framework as a recognizable subclass. The new framework does not, in general, remain confined to the old subclass under its own dynamics. The old theory survives. It does not remain sufficient.

Remark 1.1 (On the word “collapse”). Throughout this book, *collapse* is a technical term. Collapse does not mean falsification, inconsistency, or catastrophic failure. It means precisely the loss of recovery while embedding is retained. The old theory does not become false. Its theorems remain valid within the embedded subclass. What fails is the assumption that the original object class is sufficient for all relevant purposes. Collapse is the formal event at which the old ontology becomes a regime rather than a universe.

1.2 Continuation as Primitive

Before introducing the formal apparatus, it is worth naming the conceptual reorientation that the framework requires.

Traditional foundations of mathematics begin with objects: sets, points, functions, distributions. Objects are the primitive substrate; relations and operations are defined over them. The question is always: what are the things, and what can we say about them?

The framework developed in this book begins differently. The primitive question is not what a system is but whether an admissible continuation of it exists. Given a state of a system, the primary datum is whether the governing operations can be applied to produce a valid successor state within the current ontology.

Collapse is the event at which admissible continuation disappears. Enlargement is the event at which admissible continuation is restored through a larger ontology.

This reorientation has a formal expression in the admissible successor space \mathcal{A}_B introduced below. It has a philosophical expression in the preface's observation that objects appear fundamental only because they survive the dynamics used to generate them. An ontology is not a collection of primitive things. It is a collection of states that admit continuation under the governing operations.

1.3 Closure Operators and Dynamic Closure

Definition 1.2 (Dynamical Ontology). A *dynamical ontology* is a pair $(\mathcal{O}, \mathcal{F})$ where \mathcal{O} is a class of objects and $\mathcal{F} = \{F_\alpha\}_{\alpha \in A}$ is

a family of governing operations, each $F_\alpha : \mathcal{O} \rightarrow \mathcal{O}$ (possibly partial).

Definition 1.3 (Dynamic Closure). The *closure* of a class $S \subseteq \mathcal{O}$ under \mathcal{F} is

$$\text{cl}_{\mathcal{F}(S)} = \bigcup_{n=0}^{\infty} \mathcal{F}^{(n)}(S),$$

where $\mathcal{F}^{(0)}(S) = S$ and $\mathcal{F}^{(n+1)}(S) = \bigcup_{\alpha} F_\alpha(\mathcal{F}^{(n)}(S))$. The class S is *dynamically closed* under \mathcal{F} if $\text{cl}_{\mathcal{F}(S)} = S$.

Proposition 1.4. S is dynamically closed under \mathcal{F} if and only if $F_\alpha(S) \subseteq S$ for all $\alpha \in A$.

Proof. If $F_\alpha(S) \subseteq S$ for all α , then by induction $\mathcal{F}^{(n)}(S) \subseteq S$ for all n , so $\text{cl}_{\mathcal{F}(S)} \subseteq S$. Since $S = \mathcal{F}^{(0)}(S) \subseteq \text{cl}_{\mathcal{F}(S)}$, equality holds. Conversely, if $\text{cl}_{\mathcal{F}(S)} = S$ then $\mathcal{F}^{(1)}(S) \subseteq S$. \square

Example 1.5 (Smooth surfaces under area minimization). Let \mathcal{O} be smooth compact hypersurfaces with fixed boundary. In dimensions $n \leq 7$, area-minimizing competitors of smooth surfaces are smooth: \mathcal{O} is dynamically closed. In dimension $n = 8$, the Simons cone is an area-minimizing object that is not smooth: the operation escapes \mathcal{O} . Dynamic closure fails.

Example 1.6 (Real numbers under complex dynamics). Let $\mathcal{O} = \mathbb{R} \subset \mathbb{C} = \mathcal{O}'$ and $F(z) = z^2 + i$. For any $x \in \mathbb{R}$, $F(x) = x^2 + i \notin \mathbb{R}$. Embedding holds; dynamic closure fails immediately.

1.4 Embedding and Recovery

Definition 1.7 (Ontological Embedding). A *ontological embedding* is a structure-preserving injection $\iota : \mathcal{O} \hookrightarrow \mathcal{O}'$ between dynamical ontologies $(\mathcal{O}, \mathcal{F})$ and $(\mathcal{O}', \mathcal{F}')$ with $\mathcal{F} \subseteq \mathcal{F}'$.

Definition 1.8 (Ontological Recovery). Given $\iota : \mathcal{O} \hookrightarrow \mathcal{O}'$, the original ontology is *recoverable* within $(\mathcal{O}', \mathcal{F}')$ if

$$\mathcal{F}'(\mathcal{O}') \subseteq \iota(\mathcal{O}).$$

Lemma 1.9. *Recovery implies that $\iota(\mathcal{O})$ is dynamically closed under \mathcal{F}' .*

Proof. If $\mathcal{F}'(\mathcal{O}') \subseteq \iota(\mathcal{O})$, then in particular $\mathcal{F}'(\iota(\mathcal{O})) \subseteq \mathcal{F}'(\mathcal{O}') \subseteq \iota(\mathcal{O})$. \square

Proposition 1.10 (Embedding Does Not Imply Recovery). *There exist dynamical ontologies and embeddings for which $\iota(\mathcal{O})$ is not dynamically closed under \mathcal{F}' .*

Proof. Example 1.6: $\mathbb{R} \hookrightarrow \mathbb{C}$ under $F(z) = z^2 + i$. For $x \in \mathbb{R}$, $F(x) = x^2 + i \notin \mathbb{R}$. \square

1.5 Strong and Weak Survival

Definition 1.11 (Weak and Strong Survival). After enlargement $\iota : \mathcal{O} \hookrightarrow \mathcal{O}'$:

- *Weak survival* holds if $\mathcal{O} \hookrightarrow \mathcal{O}'$.

- *Strong survival* holds if weak survival holds and every theorem T valid universally in \mathcal{O} remains valid universally in \mathcal{O}' : $\mathcal{O} \vDash T \Rightarrow \mathcal{O}' \vDash T$.

Theorem 1.12. *Strong survival implies weak survival. The converse fails.*

Proof. Forward: immediate. Converse: smooth area-minimizing hypersurfaces embed in rectifiable currents (weak survival), but the Simons cone is a current that is not smooth (strong survival fails for the theorem “area-minimizing objects are smooth”). \square

Definition 1.13 (Regime Theorem). A theorem T that holds universally in \mathcal{O} but only under additional conditions in \mathcal{O}' is a *regime theorem*. The conditions under which it holds form the *recoverable regime* for T .

Ontologies as Residues

The objects of a theory appear fundamental only because they survive the dynamics used to generate them. Smooth surfaces appear to be the natural objects of variational geometry because they survive area minimization in low dimensions. Classical probability distributions appear to be the natural objects of inference because they survive Bayesian conditioning in commutative settings. Semantic scopes appear to be the natural objects of meaning because they survive admissible resolution in well-structured discourse.

In each case, the apparent fundamentality of the object class is a feature of the recoverable regime, not of the world. The objects are what persist. They are residues of the dynamics, not its foundations.

This observation inverts the usual direction of mathematical ontology. Traditionally, one begins with objects and derives dynamics. The framework developed here suggests the opposite order: dynamics come first, and objects are what they leave invariant.

1.6 Invariant Systems and Collapse

Definition 1.14 (Invariant System). An *invariant system* $\mathcal{V} = \{V_1, \dots, V_k\}$ for $(\mathcal{O}, \mathcal{F})$ is a collection of properties jointly required of all admissible objects.

Definition 1.15 (Admissible Successor Space and Inadmissibility Pressure).

$$\mathcal{A}_B = \{B' \in \mathcal{O} \mid B' \text{ satisfies all } V_i \in \mathcal{V}\}.$$

$$\mathcal{J}(B, F) = \text{dist}(F(B), \mathcal{A}_B).$$

The admissible successor space is defined by the invariant system, not by which operations succeed. This prevents circularity.

Definition 1.16 (Refuse and Collapse). An operation F is *refused* at B if $\mathcal{J}(B, F) \geq \theta > 0$ while $\mathcal{A}_B \neq \emptyset$: local failure, global ontology intact.

The ontology *collapses* when $\mathcal{A}_B = \emptyset$: no object in \mathcal{O} satisfies the full invariant system. Admissible continuation has ceased to exist within the current ontology.

1.7 Obstruction Theory

Definition 1.17 (Obstruction Set).

$$\mathcal{W}(\mathcal{O}, \mathcal{O}', \mathcal{F}') = \{w \in \mathcal{O}' \setminus \iota(\mathcal{O}) \mid w \in \mathcal{F}'(\mathcal{O}')\}.$$

An element $w \in \mathcal{W}$ is an *obstruction*: an object produced by the governing operations that belongs to the enlarged ontology but not to the original one.

Theorem 1.18 (Obstruction Implies Non-Recovery). *If $\mathcal{W} \neq \emptyset$ then recovery fails: $\mathcal{F}'(\mathcal{O}') \not\subseteq \iota(\mathcal{O})$.*

Proof. Take $w \in \mathcal{W}$. Then $w \in \mathcal{F}'(\mathcal{O}')$ and $w \notin \iota(\mathcal{O})$. □

Remark 1.19 (Obstruction versus Collapse). Theorem 1.18 establishes non-recovery from an obstruction. Non-recovery and collapse are related but distinct. Collapse ($\mathcal{A}_B = \emptyset$) concerns the original ontology's inability to provide admissible continuations. Non-recovery concerns the enlarged ontology's inability to remain confined to the original. In all three main case studies, the invariant system includes dynamic closure as a requirement, so the two conditions coincide. We will make this explicit in each case.

Theorem 1.20 (Irreversibility of Enlargement). *Suppose $\iota : \mathcal{O} \hookrightarrow \mathcal{O}'$, recovery fails, and $\text{cl}_{\mathcal{F}'}(\iota(\mathcal{O})) = \mathcal{O}'$ (the enlarged ontology is reachable from the embedded subclass). Then $\iota(\mathcal{O})$ is not dynamically closed under \mathcal{F}' .*

Proof. Since recovery fails, $\mathcal{F}'(\mathcal{O}') \not\subseteq \iota(\mathcal{O})$. By reachability, $\mathcal{O}' = \text{cl}_{\mathcal{F}'}(\iota(\mathcal{O}))$, so there exists a finite chain $x_0 \in \iota(\mathcal{O})$, $x_1 = F'_{\alpha_1}(x_0)$, ..., $x_n = F'_{\alpha_n}(x_{n-1})$ with $x_n \notin \iota(\mathcal{O})$ for the first such n . Then $F'_{\alpha_n}(x_{n-1}) \notin \iota(\mathcal{O})$ while $x_{n-1} \in \iota(\mathcal{O})$: $\iota(\mathcal{O})$ is not closed under F'_{α_n} . \square

Definition 1.21 (Obstruction Accessibility). An obstruction $w \in \mathcal{W}$ is *accessible* if there exists a feasible procedure P such that P produces w or a witness to w 's existence. Accessibility is relative to available tools and changes as methods develop.

1.8 Closure as Reflection: A Categorical View

The framework has a natural expression in category theory.

Definition 1.22 (Reflective Subcategory). A subcategory $\mathbf{O} \subset \mathbf{O}'$ is *reflective* if the inclusion functor $\iota : \mathbf{O} \hookrightarrow \mathbf{O}'$ has a left adjoint $R : \mathbf{O}' \rightarrow \mathbf{O}$, called the *reflector*. The adjunction $R \dashv \iota$ gives natural isomorphisms $\text{Hom}_{\mathbf{O}}(R(X), A) \cong \text{Hom}_{\mathbf{O}'}(X, \iota(A))$ for all $X \in \mathbf{O}'$, $A \in \mathbf{O}$.

Reflective subcategories are precisely those that are “closed under the operations of the larger category in a recoverable

way.” The reflector provides a canonical projection back into the small category.

Proposition 1.23 (Recovery as Reflection). *Ontological recovery holds if and only if the inclusion $\iota : \mathbf{O} \hookrightarrow \mathbf{O}'$ participates in a reflective adjunction $R \dashv \iota$ such that R commutes with the governing operations: $R \circ F' \cong F \circ R$.*

Corollary 1.24 (Collapse as Failure of Reflection). *Ontological collapse occurs when no reflector $R : \mathbf{O}' \rightarrow \mathbf{O}$ compatible with the governing operations exists. There is no canonical way to project outputs of \mathcal{F}' back into \mathbf{O} .*

Example 1.25 (GMT in Categorical Terms). Let \mathbf{O} be smooth minimal hypersurfaces and \mathbf{O}' integer-multiplicity rectifiable currents. The inclusion does not have a reflector compatible with area minimization in dimension eight: there is no canonical way to smooth out the Simons cone while preserving its minimizing property. Reflection fails; collapse is confirmed.

1.9 The Six-Stage Pattern

Stage 1. Operational Failure (Refuse). $\mathcal{J}(B, F) \geq \theta$ while $\mathcal{A}_B \neq \emptyset$. Operation blocked locally; ontology intact globally.

Stage 2. Ontological Collapse. $\mathcal{A}_B = \emptyset$. Admissible continuation has ceased to exist within the current ontology. This is an objective mathematical fact, independent of recognition.

Stage 3. Enlargement Availability. $\mathcal{O}' \supset \iota(\mathcal{O})$ is constructed with $\mathcal{F}'(\mathcal{O}') \subseteq \mathcal{O}'$. May precede recognized collapse.

Stage 4. Obstruction Existence. $\mathcal{W} \neq \emptyset$. A specific object certifies that recovery fails.

Stage 5. Obstruction Accessibility. A feasible procedure produces a witness. The gap between Stages 4 and 5 is the *accessibility gap*, governed by the structure of the obstruction region relative to available tools — not by the sociology of the field.

Stage 6. Ontological Commitment. The field accepts the enlarged ontology as necessary. Universal claims of the old theory are reinterpreted as regime theorems.

Meta-Theorem 1.26 (Closure-Recovery Principle). *Let $(\mathcal{O}, \mathcal{F})$ be a dynamical ontology. If $\text{cl}_{\mathcal{F}(\mathcal{O})} \neq \mathcal{O}$, then either \mathcal{F} must be restricted or \mathcal{O} must be enlarged. Historically, enlargement is the repeated outcome.*

Part II

The Geometry Case: A Completed Collapse

Chapter 2

Functionals and the First Enlargement

Bernoulli asked for the fastest curve.

Lagrange realized the question was about all curves at once.

2.1 The Brachistochrone and the Birth of Functionals

In 1696 Johann Bernoulli posed a problem that looks, at first, like an exercise in mechanics. A bead slides without friction along a wire connecting two points in a vertical plane. What shape should the wire take so that the bead reaches the lower point in the shortest possible time?

The answer is a cycloid. But Bernoulli understood that the significant achievement was not the answer but the method. To find the fastest wire, one does not examine a particular wire. One examines all wires simultaneously and asks which

is stationary with respect to travel time. Travel time is not a function of a real variable; it is a function of an entire curve: a *functional*. The problem of finding curves making a functional stationary is the subject of the calculus of variations.

Let the wire be described by $y(x)$ with $y(x_0) = y_0, y(x_1) = y_1$. Conservation of energy gives speed $v = \sqrt{2g(y_0 - y)}$ at height y . The travel time is

$$T[y] = \int_{x_0}^{x_1} \frac{\sqrt{1 + (y')^2}}{\sqrt{2g(y_0 - y)}} dx.$$

This is a functional: a map from the space of curves to the real numbers.

2.2 The Euler-Lagrange Equation

Theorem 2.1 (Euler-Lagrange). *If y is C^2 , satisfies fixed boundary conditions $y(a) = A, y(b) = B$, and $\delta J[y] = 0$ for all compactly supported smooth perturbations η vanishing at the endpoints, where $J[y] = \int_a^b L(x, y, y') dx$, then*

$$\frac{\partial L}{\partial y} - \frac{d}{dx} \frac{\partial L}{\partial y'} = 0.$$

Proof. Set $y_\varepsilon = y + \varepsilon\eta$. Differentiating at $\varepsilon = 0$:

$$\left. \frac{d}{d\varepsilon} \right|_0 J[y_\varepsilon] = \int_a^b \left(\frac{\partial L}{\partial y} \eta + \frac{\partial L}{\partial y'} \eta' \right) dx = 0.$$

Integrating the second term by parts and using $\eta(a) = \eta(b) =$

0:

$$\int_a^b \left(\frac{\partial L}{\partial y} - \frac{d}{dx} \frac{\partial L}{\partial y'} \right) \eta \, dx = 0.$$

Since η is arbitrary, the integrand vanishes everywhere. \square

Remark 2.2 (The first ontological enlargement). The passage from points to functionals is itself an ontological enlargement. The brachistochrone could not be stated without the concept of a functional. Functions on curves embed in the space of all functions on an interval, but the variational calculus operates on function spaces whose governing dynamics do not remain confined to finite-dimensional subspaces. Embedding holds; recovery fails immediately. The calculus of variations is the enlarged ontology that restored closure, at the cost of requiring the Euler-Lagrange equation rather than simple differentiation.

2.3 The Plateau Problem and Minimal Surfaces

For a surface expressed as a graph $z = f(x, y)$, the minimal surface equation (vanishing mean curvature $H = 0$) is

$$(1 + f_y^2)f_{xx} - 2f_x f_y f_{xy} + (1 + f_x^2)f_{yy} = 0.$$

Physical soap films satisfy this equation: area minimization produces vanishing mean curvature. The Plateau problem asks whether, given any closed curve, a minimal surface

spanning it exists.

2.4 Bernstein's Theorem and the Illusion of Universality

Theorem 2.3 (Bernstein, 1915 [Ber15]). *Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be C^2 with minimal graph. Then f is affine: the only complete minimal graphs in \mathbb{R}^3 are planes.*

This was extended through dimension seven by De Giorgi [De 65], Almgren [Alm66], and Simons [Sim68]:

Theorem 2.4 (Bernstein in Low Dimensions). *The Bernstein theorem holds for entire minimal hypergraphs $f : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ for $n \leq 8$: any C^2 solution of the minimal surface equation defined on all of \mathbb{R}^{n-1} must be linear.*

With each dimensional extension the expectation hardened: minimization enforces simplicity universally. The invariant system (smooth, minimizing, globally defined) appeared jointly satisfiable in all dimensions. The collapse was present in dimension eight long before it was recognized.

Chapter 3

The Simons Cone and Obstruction

The cone is not important because it is singular.

It is important because it is necessary.

3.1 The Construction

In \mathbb{R}^8 define

$$\mathcal{C} = \{(x_1, \dots, x_8) \in \mathbb{R}^8 \mid x_1^2 + x_2^2 + x_3^2 + x_4^2 = x_5^2 + x_6^2 + x_7^2 + x_8^2\}.$$

This is the *Simons cone* [Sim68]: a seven-dimensional cone with vertex at the origin, smooth everywhere except the origin.

Theorem 3.1 (Simons 1968; Bombieri-De Giorgi-Giusti 1969 [Sim68; BDG69]). *The Simons cone $\mathcal{C} \subset \mathbb{R}^8$ is area-minimizing. Moreover, the Bernstein theorem fails for $n \geq 9$: there exist complete minimal hypergraphs in \mathbb{R}^n that are not affine hyperplanes.*

Dimension eight is the sharp boundary.

3.2 Dimension as Control Parameter

The mechanism is competition between curvature cost and area gain. In low dimensions, singularities are energetically penalized: curvature concentrates at a singular point at greater cost than the area saved. Additional dimensions distribute curvature across more principal directions, weakening the penalty. At dimension eight, the balance tips.

Stability is established by the second variation:

Definition 3.2 (Second Variation and Stability). For smooth minimal Σ and compactly supported ϕ :

$$\delta^2 \mathcal{A}(\Sigma)[\phi, \phi] = \int_{\Sigma} (|\nabla \phi|^2 - |A|^2 \phi^2) d\mathcal{H}^{n-1}.$$

Σ is *stable* if $\delta^2 \mathcal{A}[\phi, \phi] \geq 0$ for all compactly supported ϕ .

Simons showed that \mathcal{C} satisfies the estimates needed for stability in \mathbb{R}^8 but not in lower-dimensional analogues. The balance between $|\nabla \phi|^2$ and $|A|^2 \phi^2$ tips at precisely the ambient dimension where the cone lives.

3.3 The Cone as Obstruction

In the language of Chapter 1, identifying each component:

Original ontology \mathcal{O} : smooth minimal hypersurfaces. *Governing operation*: area minimization. *Invariant system*: V_1 smoothness, V_2 area-minimization, V_3 global existence for all

boundary data. *Enlarged ontology* \mathcal{O}' : integer-multiplicity rectifiable currents. *Obstruction*: the Simons cone $\mathcal{C} \in \mathcal{O}' \setminus \iota(\mathcal{O})$, produced by area minimization, certifying that recovery fails.

By Theorem 1.18, the cone certifies non-recovery. Since the invariant system includes dynamic closure, non-recovery and collapse coincide: the smooth-surface ontology has no admissible continuation for the cone's boundary conditions.

The cone does not create the failure. It certifies a collapse already present in the mathematics of \mathbb{R}^8 . The cone is what makes the collapse impossible to ignore.

3.4 Smoothness Becomes a Regime

After Theorem 3.1, “area-minimizing surfaces are smooth” is reinterpreted as a regime theorem. It holds in dimensions ≤ 7 and as a statement about the regular part of minimizing currents in higher dimensions.

Theorem 3.3 (Regularity, Sharp Form [Fed69; Sim83]). *For an area-minimizing integer-multiplicity rectifiable n -current in \mathbb{R}^{n+1} :*

- (1) *The regular set $\text{Reg}(T)$ is open and dense in $\text{supp}(T)$.*
- (2) $\dim_{\mathcal{H}} \text{Sing}(T) \leq n - 7$.
- (3) *In dimensions $n \leq 6$: $\text{Sing}(T) = \emptyset$.*
- (4) *In dimension $n = 7$: $\text{Sing}(T)$ consists of isolated points.*

The number seven appears because it is the lowest dimension in which the Simons cone (a dimension-seven object)

can be a stable minimizer. Below the threshold, singular minimizers cannot persist; above it, they can.

Chapter 4

Geometric Measure Theory as Enlarged Ontology

Classical differential geometry assumes smooth surfaces.

Geometric measure theory asks what minimization actually produces.

4.1 The Closure Failure of Smooth Surfaces

Minimizing sequences of smooth surfaces need not have convergent subsequences in the smooth category. Curvatures can concentrate; singularities can develop; topological type can change. The class of smooth surfaces is not closed under sequential limits, which variational methods require.

4.2 Currents: Closure Restored

Definition 4.1 (Current and Mass [Fed69]). A k -dimensional current in \mathbb{R}^n is a continuous linear functional $T : \Omega_c^k(\mathbb{R}^n) \rightarrow \mathbb{R}$, with boundary $\partial T(\omega) = T(d\omega)$ and mass $\mathbf{M}(T) = \sup\{T(\omega) : \sup|\omega| \leq 1\}$.

An integer-multiplicity rectifiable k -current is $T(\omega) = \int_M \langle \omega, \xi \rangle \theta d\mathcal{H}^k$ where M is k -rectifiable, ξ is an orientation, and $\theta : M \rightarrow \mathbb{Z}_{>0}$.

Smooth oriented surfaces are the special case $\theta \equiv 1$. The Simons cone is a rectifiable current with a singularity at the origin. Both live in the same enlarged ontology.

Theorem 4.2 (Federer-Fleming Compactness [Fed69]). A sequence of integer-multiplicity rectifiable k -currents with uniformly bounded mass and boundary mass has a weakly convergent subsequence to an integer-multiplicity rectifiable k -current.

Remark 4.3 (Closure restored). Theorem 4.2 is precisely the statement that the enlarged ontology is dynamically closed under the sequential limits that variational methods require. Minimizing sequences that escape the smooth category are contained within the current category. The governing operation closes on the enlarged ontology.

Theorem 4.4 (Existence of Area-Minimizing Currents). For any rectifiable $(k - 1)$ -cycle Γ , there exists an integer-multiplicity rectifiable k -current T with $\partial T = \Gamma$ achieving the minimum mass.

4.3 Tangent Cones and Local Models

Definition 4.5 (Tangent Cone). Let T be a current, $x \in \text{Sing}(T)$. The *tangent cone* at x is the weak limit as $r \rightarrow 0$ of the rescaled currents $T_{x,r}(\omega) = r^{-n}T((\tau_x \circ \delta_{1/r})\# \omega)$.

Tangent cones are local models for singularities. The Simons cone appears as the tangent cone for a generic class of singularities in dimension eight. It is an attractor in the space of rescaled currents: a renormalization fixed point. This explains why it is not merely a counterexample but a structural object. Singularities, in the vicinity of this particular renormalization fixed point, look like cones.

Return and Recovery

The purpose of enlargement is not escape from the original ontology. It is the recovery of those conditions under which the original ontology remains valid.

Regularity theory is a theory of return. Given a minimizing current in the enlarged ontology, regularity theory determines where and to what extent the original smooth description is recoverable. The regular set $\text{Reg}(T)$ is the region of return: where the object behaves as the old ontology predicted. The singular set $\text{Sing}(T)$ is where return is impossible.

This structure repeats across all the case studies in this book. Asymptotic frequentism in quantum inference is a theorem of return: the posterior eventually reaches the true state, even if the path is non-monotone. Scope

completion in Spheredrop is a theorem of return: the load eventually reaches zero, even if intermediate operations temporarily increase it.

Return to the old ontology is possible in the generic case. The enlarged ontology does not eliminate the old one; it situates it. It explains when the old description holds and when it does not. The enlarged framework is valuable precisely because it provides a theory of return.

4.4 The GMT Case: Five Claims

The GMT case establishes five claims that the subsequent case studies reproduce in modified form.

Collapse is objective and precedes recognition. The invariant system was unsatisfiable in dimension eight before the Simons cone was constructed.

The enlarged ontology must be available before an obstruction can be recognized. GMT existed before the cone.

The accessibility gap is structural, not sociological. The cone was constructible with available tools within approximately a decade.

Strong survival fails while weak survival holds. The universal smoothness theorem becomes a regime theorem.

The enlarged ontology has its own structure, richer than the original. Regularity theory, tangent cones, mean curvature flow: these research programs do not exist within the smooth category. The enlargement opens new territory.

Part III

The Inference Case: An Ongoing Collapse

Chapter 5

Classical Bayesian Inference and Its Geometry

Probability is not a property of the world. It is a property of our relationship to the world. But that relationship has a geometry.

5.1 Inference as Dynamical Process

The original ontology is the class of probability distributions $\mathcal{O} = \{\pi : \Theta \rightarrow \mathbb{R}_{\geq 0} \mid \int \pi = 1\}$ over a hypothesis space Θ . The governing operation is Bayesian conditioning: $F_e(\pi)(h) = \pi(h \mid e) = \pi(e \mid h)\pi(h)/\pi(e)$. The invariant system consists of V_1 (classical joint state existence), V_2 (unique conditioning), and V_3 (monotone improvement under repetition). All three are jointly satisfiable in the commutative setting.

5.2 Kullback-Leibler Divergence and the Fisher Metric

Definition 5.1 (Kullback-Leibler Divergence). $D_{\text{KL}}(p\|q) = \int p(x) \log(p(x)/q(x)) dx \geq 0$, with equality iff $p = q$ almost everywhere.

Definition 5.2 (Fisher Information Metric). $g_{ij}(\theta) = \mathbb{E}_{p_\theta}[(\partial_i \log p_\theta)(\partial_j \log p_\theta)]$

The Fisher metric is, up to rescaling, the unique Riemannian metric on statistical manifolds invariant under sufficient statistics [Che82]. Bayesian conditioning is an information projection in this geometry: the posterior minimizes relative entropy to the prior subject to matching the expected sufficient statistics of the evidence.

5.3 The Data Processing Inequality and No Epistemic Harm

Theorem 5.3 (Data Processing Inequality). *For any measurable $\phi : \mathcal{X} \rightarrow \mathcal{Y}$ and distributions p, q : $D_{\text{KL}}(\phi_*p\|\phi_*q) \leq D_{\text{KL}}(p\|q)$.*

Sketch. Chain rule: $D_{\text{KL}}(p_{XY}\|q_{XY}) = D_{\text{KL}}(p_X\|q_X) + D_{\text{KL}}(p_{Y|X}\|q_{Y|X})$. Processing by ϕ marginalizes the non-negative second term.

□

Corollary 5.4 (No Epistemic Harm, Classical). *Let π^* be the true distribution, π_n the posterior after n observations. Then*

$D_{\text{KL}}(\pi^* \|\pi_{n+1}) \leq D_{\text{KL}}(\pi^* \|\pi_n)$. Repeated Bayesian updating weakly decreases divergence from truth.

This is invariant V_3 : monotone improvement under the governing operation.

5.4 Asymptotic Frequentism

Theorem 5.5 (Bernstein-von Mises). *Under mild regularity conditions, the posterior $\pi_n(\cdot \mid x_1, \dots, x_n)$ converges in total variation to δ_{θ^*} as $n \rightarrow \infty$, at a rate asymptotically independent of the prior π_0 .*

The true distribution is the attractor; the prior is the initial condition. In the classical case, no epistemic harm (local monotonicity) and asymptotic frequentism (global convergence) both hold. In the quantum case, they may come apart.

5.5 Schrödinger Bridges and Process Updating

Definition 5.6 (Schrödinger Bridge [Sch32]). Let \mathbb{P} be Wiener measure and μ_0, μ_1 measures on \mathbb{R}^d . The *Schrödinger bridge* is

$$\mathbb{Q}^* = \operatorname{argmin}\{D_{\text{KL}}(\mathbb{Q} \|\mathbb{P}) \mid \mathbb{Q}_0 = \mu_0, \mathbb{Q}_1 = \mu_1\}.$$

The Sinkhorn algorithm [Sin67] computes Schrödinger bridges by alternately projecting onto the marginal constraints.

This is inference at the process level: instead of updating a belief state, one updates a coupling. The enlarged ontology of processes embeds the original ontology of states; whether recovery holds is the question the quantum case approaches.

Chapter 6

Quantum Inference and the Prospective Collapse

Quantum mechanics is not probability theory with complex numbers.
It is probability theory with a different geometry.

6.1 Density Matrices and Noncommutativity

Definition 6.1 (Density Matrix and Quantum Relative Entropy). A *density matrix* on \mathcal{H} is $\rho \in \mathcal{L}(\mathcal{H})$ with $\rho \geq 0$ and $\text{tr}(\rho) = 1$. For density matrices with $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$:

$$D(\rho\|\sigma) = \text{tr}(\rho \log \rho - \rho \log \sigma) \geq 0.$$

The fundamental structural difference: density matrices generically fail to commute, $[\rho, \sigma] \neq 0$. This destroys the

canonical joint-state structure classical inference requires.

Theorem 6.2 (Quantum Data Processing Inequality [WWY15]).

For any CPTP map \mathcal{E} : $D(\mathcal{E}(\rho)\|\mathcal{E}(\sigma)) \leq D(\rho\|\sigma)$.

6.2 The Petz Recovery Map

Definition 6.3 (Petz Recovery Map [Pet88]). Let $\sigma \in \mathcal{D}(\mathcal{H}_{AB})$ with marginal $\sigma_B = \text{tr}_A(\sigma)$. The *Petz recovery map* is

$$\mathcal{R}_{\sigma, \text{tr}_A}(X) = \sigma^{1/2}(\sigma_B^{-1/2} X \sigma_B^{-1/2} \otimes \mathbf{1}_A) \sigma^{1/2}.$$

Proposition 6.4 (Properties of the Petz Map). (1) $\mathcal{R}_{\sigma, \text{tr}_A}$ is CPTP.

(2) *Quantum detailed balance*: $D(\rho\|\sigma) = D(\text{tr}_A(\rho)\|\sigma_B) + D(\rho\|\mathcal{R}_{\sigma, \text{tr}_A}(\text{tr}_A(\rho)))$.

(3) *In the commutative regime, it reduces to classical Bayesian conditioning.*

Property (3) confirms $\mathcal{O} \hookrightarrow \mathcal{O}'$: classical inference embeds in quantum inference. Property (2) is the quantum Bayes rule. The question is recovery.

6.3 A Worked Petz Recovery Example

Let

$$\rho = \begin{pmatrix} 0.7 & 0.25 \\ 0.25 & 0.3 \end{pmatrix}, \quad \sigma = \begin{pmatrix} 0.6 & 0 \\ 0 & 0.4 \end{pmatrix}.$$

These matrices do not commute: $[\rho, \sigma] \neq 0$, placing the system outside the classical regime. Consider the dephasing

channel $\mathcal{E}(X) = \text{diag}(X_{11}, X_{22})$:

$$\mathcal{E}(\rho) = \begin{pmatrix} 0.7 & 0 \\ 0 & 0.3 \end{pmatrix}, \quad \mathcal{E}(\sigma) = \sigma.$$

The quantum data-processing inequality guarantees $D(\mathcal{E}(\rho)\|\mathcal{E}(\sigma)) \leq D(\rho\|\sigma)$. Information is lost under coarse-graining, as expected.

The Petz recovery map produces a recovered state $\rho_1 = \mathcal{R}_{\sigma, \mathcal{E}}(\mathcal{E}(\rho))$. The diagnostic quantity is

$$\Delta_n = D(\rho^*\|\rho_{n+1}) - D(\rho^*\|\rho_n),$$

where ρ^* is the target state and ρ_n the state after n Petz recovery cycles.

Classical Bayesian inference predicts $\Delta_n \leq 0$ for every n (Corollary 5.4). The conjectural obstruction regime predicts the existence of specific triples $(\rho, \sigma, \mathcal{E})$ and targets ρ^* for which $\Delta_n > 0$ for some finite n , even though $D(\rho^*\|\rho_n) \rightarrow 0$ as $n \rightarrow \infty$.

Such a trajectory temporarily moves away from the target before ultimately converging. This is not a failure of inference. It is a failure of monotonicity: the failure of V_3 while global convergence (the attractor structure) may survive. The obstruction is not non-convergence. The obstruction is the failure of monotonicity itself, which would certify that the classical invariant system is unsatisfiable in the noncommutative regime.

The GMT parallel is exact. Minimizing currents exist and

generically converge to area-minimizing limits (global existence survives); but not every minimizing current is smooth (local regularity fails). Quantum inference may generically converge to the true state (attractor structure survives); but not every inference trajectory is monotone (local monotonicity fails).

6.4 The Invariant System and What Collapses

The classical invariant system for quantum inference:

V_1 : a classical joint state encodes the prior-evidence relationship.

V_2 : conditioning is uniquely determined by the joint state.

V_3 : $D(\rho^* \parallel \rho_{n+1}) \leq D(\rho^* \parallel \rho_n)$ for all n .

Under noncommutativity, V_1 and V_2 are structurally unsatisfiable: quantum systems generically lack classical joint states. Whether V_3 fails is the open question that determines whether the prospective collapse is established.

Conjecture 6.5 (Noncommutativity Threshold). *There exists a measure η of effective noncommutativity and a critical value η_c such that: for $\eta < \eta_c$, the Petz map satisfies V_3 ; for $\eta > \eta_c$, violations of V_3 are generic; the transition is sharp.*

This is a falsifiable prediction. Gradual degradation of monotonicity across all η would weaken the Simons cone analogy considerably. Sharp threshold behavior would confirm the structural parallel quantitatively.

6.5 Established versus Prospective Collapse

Definition 6.6 (Established and Prospective Collapse). A collapse is *established* if there is a formal proof that $\mathcal{A}_B = \emptyset$ in \mathcal{O} for the relevant invariant system. A collapse is *prospective* if strong structural and numerical evidence exists but formal proof is absent.

The GMT case is established. The quantum Bayes case is prospective. This distinction does not weaken the framework; it makes it diagnostic rather than merely retrospective. A framework that only classifies completed revolutions is a catalogue.

Part IV

The Formal Case: Spherepop as Abstract Model

Chapter 7

Scope Graphs and Semantic Topology

A bubble is not a thing. It is a commitment deferred.

7.1 The Spherepop Calculus: Definitions

Definition 7.1 (Semantic Bubble). A *semantic bubble* is a triple $B = (C, E, U)$ where C is a contextual binding set, E is an expectation structure, and $U(B) \geq 0$ is the unresolved semantic load. B is *open* if $U(B) > 0$, *closed* if $U(B) = 0$.

Definition 7.2 (Containment Structure, Scope Stack, Load Functional). A *containment structure* is (\mathcal{B}, \prec) with \prec a strict partial order. The *scope stack* is $\Sigma = [B_1 \prec \dots \prec B_n]$. The *semantic load* is $L(\Sigma) = \sum_i w_i U(B_i)$ where $w_i = f(d_i, s_i, r_i, c_i)$ weights depth, salience, recency, and causal centrality.

7.2 Scope Graphs

Containment trees are too restrictive for cross-cutting dependencies. Human reasoning, mathematical proof, software architecture, and scientific explanation frequently involve structures where a scope depends on another without being spatially contained within it.

Definition 7.3 (Scope Graph). A *scope graph* is a finite directed graph $G = (V, E)$ with semantic load $U : V \rightarrow \mathbb{R}_{\geq 0}$. Vertices are semantic bubbles; edge $B_i \rightarrow B_j$ means admissibility of B_i depends on resolution of B_j .

Containment trees appear as scope graphs with at most one outgoing dependency per vertex and no directed cycles. The scope graph separates logical dependence from spatial containment.

Definition 7.4 (Dependency Closure). The *dependency closure* $\text{cl}(G)$ is the smallest graph containing G and closed under transitivity: $(B_i, B_j) \in \text{cl}(G)$ whenever there exists a directed path from B_i to B_j .

Proposition 7.5. $\text{cl}(\text{cl}(G)) = \text{cl}(G)$.

Proof. The first closure adjoins every transitive dependency. A second application introduces no additional edges. \square

Dependency closure is the Spherepop analogue of the topological closure operator: it closes a semantic structure under dependency implication rather than a set under limits.

7.3 Collapse Cycles as Obstructions

Definition 7.6 (Collapse Cycle). A *collapse cycle* is a directed cycle $B_1 \rightarrow B_2 \rightarrow \dots \rightarrow B_n \rightarrow B_1$ in which every edge represents an unresolved dependency.

Theorem 7.7 (Collapse Cycle Theorem). *If a scope graph contains a collapse cycle then $\mathcal{A} = \emptyset$.*

Proof. Resolution of B_1 requires prior resolution of B_2 . Resolution of B_2 requires prior resolution of B_3 . By induction, resolution of B_n requires prior resolution of B_1 . No node may be resolved first. Therefore no admissible Pop sequence exists and $\mathcal{A} = \emptyset$. \square

The collapse cycle plays the role of the Simons cone in the graph formalism. It does not create the collapse; it certifies it. A collapse cycle is a stable dependency configuration that makes the collapse impossible to ignore, exactly as the Simons cone is a stable singular configuration certifying that smooth minimizers do not exist for the cone's boundary conditions.

7.4 Minimal Obstructions

Definition 7.8 (Minimal Obstruction). A subgraph $H \subseteq G$ is a *minimal obstruction* if H satisfies Theorem 7.7 (it contains a collapse cycle) but every proper subgraph of H admits an admissible resolution sequence.

Theorem 7.9 (Existence of Minimal Obstructions). *Every finite collapsed scope graph contains a minimal obstruction.*

Proof. Let $\mathcal{O} = \{H \subseteq G : H \text{ is collapsed}\}$. This is nonempty since $G \in \mathcal{O}$. Since G is finite, \mathcal{O} has an element of minimal cardinality. Any such element is a minimal obstruction: removing any edge or vertex would restore admissibility. \square

This establishes an obstruction theory for Spherepop analogous to forbidden-minor theory in graph topology. Minimal obstructions are the smallest witnesses to collapse: the Spherepop analogues of the Simons cone.

7.5 Semantic Entropy and the Obstruction Measure

Definition 7.10 (Semantic Entropy). The *semantic entropy* of a scope graph G is

$$S(G) = \log |\mathcal{A}(G)|,$$

the log-cardinality of admissible resolution sequences ($-\infty$ when $\mathcal{A} = \emptyset$, i.e., at collapse).

High $S(G)$ means many admissible continuations: the system is far from collapse. Low $S(G)$ means few: collapse is imminent. $S(G) = -\infty$ means collapse has occurred. Semantic entropy is the Spherepop analogue of $\text{dist}(F(B), \mathcal{A}_B)$.

Theorem 7.11 (Semantic Thermodynamics). *Under pure Pop sequences, semantic entropy is non-decreasing. Under Meld and Reframe, semantic entropy may temporarily decrease before stabilizing.*

Sketch. Each admissible Pop closes one scope, removing one constraint from the parent’s admissibility condition, weakly increasing $|\mathcal{A}|$. Meld and Reframe introduce new compatibility constraints that may temporarily reduce $|\mathcal{A}|$ before they stabilize. \square

7.6 The Collapse Quotient and Irreversibility

Definition 7.12 (Collapse Quotient). Let \sim_ρ be the equivalence relation on bubble histories: two histories are equivalent when they produce the same admissible semantic representative under all available resolution sequences. The *collapsed residue* of B is B/\sim_ρ .

Construction 7.13 (Universal Property). The quotient map $q : B \rightarrow B/\sim_\rho$ satisfies: for any $f : B \rightarrow X$ constant on \sim_ρ -classes, there exists unique $\bar{f} : B/\sim_\rho \rightarrow X$ with $f = \bar{f} \circ q$.

Proposition 7.14 (Irreversibility of Spherepop Collapse). *The quotient map $q : B \rightarrow B/\sim_\rho$ does not generally admit a section. Generative history is not recoverable from the collapsed residue.*

Proof. A section would assign to each equivalence class a canonical representative history. But multiple distinct histo-

ries may produce the same residue. No choice among them is canonical without information absent from the residue. \square

Example 7.15 (Arithmetic Irreversibility). $(2 + 3)$ and $(1 + 4)$ both produce residue 5. From 5 alone, neither history is recoverable. The history is lost; the result is authoritative. This is the Spheredpop instance of the central thesis: collapse preserves embedding (both histories map to 5) but not recovery (the residue does not retract onto either history).

Theorem 7.16 (Monotonicity Under Pure Pop Sequences). *A sequence of admissible Pop operations with $w_i \geq 0$ satisfies $L(\Sigma_{k+1}) \leq L(\Sigma_k)$.*

Proof. Each admissible Pop removes bubble B_i contributing $w_i U(B_i) > 0$. The admissibility condition ensures all descendants have $U = 0$, so no new load is introduced. Hence L strictly decreases. \square

Remark 7.17 (Spheredpop predicts its own epistemic harm failure). Theorem 7.16 holds for pure Pop sequences. Meld and Reframe may temporarily increase $L(\Sigma)$. This is the Spheredpop instance of the quantum Bayes phenomenon: the simplest operations (Pop, classical conditioning) are monotone; richer operations (Meld, Reframe, quantum recovery) may violate monotonicity while preserving global convergence. The calculus predicts the pattern it is modeling.

The fundamental question of the Spherepop calculus is not: what does this bubble mean? It is: does an admissible continuation exist?

The six-stage pattern of this book is organized around the same question at every scale. The governing operations of a theory ask, at each step, whether the current state admits a valid successor. Refuse is the answer “yes, but not this operation.” Collapse is the answer “no, and none will do.”

This is why the framework is fundamentally different from accounts that emphasize complexity or abstraction as drivers of ontological change. Complexity does not force enlargement. The absence of admissible continuation forces enlargement. The Simons cone does not make minimal surface theory more complex. It makes admissible continuation impossible within the smooth category. The collapse is not a failure of understanding. It is a failure of continuation.

The new ontology is not a more complex version of the old one. It is the minimal structure in which admissible continuation is restored.

7.7 Spherepop as Formal Model

Stage	Formal Condition	Spherepop Realization
Refuse	$\mathcal{J}(B, F) \geq \theta, \mathcal{A} \neq \emptyset$	Pop with open descends
Collapse	$\mathcal{A} = \emptyset$	Collapse cycle exists
Enlargement	$\mathcal{O}' \supset \iota(\mathcal{O})$	Quotient category
Obstruction	$w \in \mathcal{O}' \setminus \iota(\mathcal{O}), w \in \mathcal{F}'(\mathcal{O}')$	Minimal obstruction su
Accessibility	Feasible $P \mapsto w$	Cascade threshold $L >$
Commitment	Enlarged ontology accepted	Quotient semantics ad

The cascade threshold L_c is the Spherepop control parameter: the analogue of dimension eight in GMT and η_c in the quantum Bayes conjecture.

Part V

Diagnostic Application: Foundation Models

Chapter 8

The Token Ontology as an Invariant System

The history of artificial intelligence may be interpreted as a search for the minimal enlarged ontology required to restore closure under the dynamics of large-scale learning.

8.1 The Token Ontology in the Framework

The token ontology expressed precisely:

Original ontology \mathcal{O} : finite token sequences over a fixed vocabulary.

Governing operation F : next-token prediction. The model learns $p_\theta(x_{n+1} \mid x_1, \dots, x_n)$, and all capabilities are understood as consequences of this distribution.

Invariant system \mathcal{V} :

V_1 (*Representation Sufficiency*): all relevant structure can be represented as token statistics.

V_2 (*Predictive Sufficiency*): next-token prediction is sufficient to explain capability formation.

V_3 (*Behavioral Sufficiency*): deployment behavior is recoverable from training behavior within the token-distribution description.

V_4 (*Interpretive Sufficiency*): internal representations require no ontology beyond token relationships to characterize model behavior and failure modes.

All four are jointly satisfiable for small models on narrow tasks. The question is whether they remain jointly satisfiable as models scale and tasks generalize.

8.2 Candidate Enlarged Ontologies

Large-scale models produce structures that behave as elements of $\mathcal{O}' \setminus \iota(\mathcal{O})$.

The *embedding ontology* takes vectors in representation space as fundamental. Geometry in that space, not token statistics, is the governing structure.

The *latent world-model ontology* takes the model's internal representations as encoding a compressed generative model of the world. The governing operation is inference within that world model.

The *mechanistic ontology* takes specific circuits, attention patterns, and feature directions as fundamental. These are

identified by interpretability research and may have no natural description as token statistics.

The *agent ontology* takes goals and beliefs as fundamental, with the governing operation being action selection.

Each embeds the token ontology: token prediction is a degenerate process in each enlarged framework. Whether any of the four invariants has become globally unsatisfiable is the open question.

8.3 Candidate Obstructions

A formal obstruction would be a capability or failure mode that cannot be expressed while preserving all four invariants simultaneously.

Systematic out-of-distribution generalization is a candidate obstruction to V_1 and V_2 : behaviors on inputs outside the training distribution that the token ontology cannot predict or explain.

Alignment failures are candidate obstructions to V_3 : systematic divergences between training and deployment behavior that require a different vocabulary to characterize.

Mechanistic features with no token-level description are candidate obstructions to V_4 : circuits or feature directions that cannot be characterized as functions of token statistics.

None is yet a formal obstruction in the sense of Theorem 1.18: no explicit construction has proven that preserving all four invariants simultaneously is impossible for any specific capability or failure mode.

8.4 Accessibility and Phenomenology

The accessibility structure differs qualitatively from the GMT and quantum Bayes cases. In GMT, the obstruction was mathematically constructible. In quantum Bayes, it may require computational exploration of large Hilbert spaces. In the foundation model case, the obstruction region may be *empirically* inaccessible: the relevant phenomena occur in deployment contexts that are difficult to study systematically.

This predicts a qualitatively different phenomenology: ontological commitment in this domain will be driven by high-profile deployment failures functioning as accessible obstructions, not by formal proofs. The engineering community will commit to an enlarged ontology before the mathematical community has established that collapse has occurred.

8.5 Framework Predictions

If the token ontology has collapsed, the framework predicts: the collapse is organized by a control parameter (scale, task complexity, or both); the transition exhibits regime structure rather than smooth degradation; the minimal enlarged ontology is whichever candidate framework is the smallest enlargement restoring closure; and the regularity theory of the enlarged ontology characterizes the conditions under which token-level descriptions remain adequate.

These are empirical predictions with specific content. They are not predictions that the token ontology will fail, but pre-

dictions about *how* a failure, if present, will be structured.

Part VI

**General Theory and
Conclusions**

Chapter 9

A General Theory of Ontological Enlargement

New objects matter because they reveal that old ontologies were not closed under their own governing dynamics.

9.1 The Central Results

Meta-Theorem 9.1 (Closure-Recovery Principle). *Let $(\mathcal{O}, \mathcal{F})$ be a dynamical ontology. If $\text{cl}_{\mathcal{F}(\mathcal{O})} \neq \mathcal{O}$, then either \mathcal{F} must be restricted or \mathcal{O} must be enlarged. Historically, enlargement is the repeated outcome.*

Theorem 9.2 (Irreversibility of Enlargement). *Suppose $\iota : \mathcal{O} \hookrightarrow \mathcal{O}'$, recovery fails, and $\text{cl}_{\mathcal{F}'}(\iota(\mathcal{O})) = \mathcal{O}'$. Then $\iota(\mathcal{O})$ is not dynamically closed under \mathcal{F}' .*

Theorem 9.3 (Collapse as Failure of Reflection). *Ontological collapse occurs when no reflector $R : \mathbf{O}' \rightarrow \mathbf{O}$ compatible with the*

governing operations exists. There is no canonical way to project outputs of \mathcal{F}' back into \mathbf{O} .

Theorem 9.4 (Regime Theorem). *Let T be a universal theorem in $(\mathcal{O}, \mathcal{V})$. After enlargement, T becomes a regime theorem: valid within a recoverable subregime $\mathcal{R} \subseteq \mathcal{O}'$ defined by additional conditions, and potentially failing outside \mathcal{R} . Characterizing \mathcal{R} precisely is the primary task of post-collapse regularity theory.*

Theorem 9.5 (Delay is Structural). *The gap between Stage 2 and Stage 6 is bounded below by the time required to identify an accessible obstruction. This bound is governed by the structure of the obstruction region relative to available tools, not by the sociology of the field.*

9.2 Comparative Table

	GMT	Quantum Bayes	Spherepop	Foundat. Models
\mathcal{O}	Smooth mini- mal hypersur- faces	Classical prob. distrib.	Well-nested stacks	Token quences
\mathcal{F}	Area mini- mization	Bayesian con- ditioning	Pop, Meld, Re- frame	Next-tok prediction
\mathcal{V}	Smooth + min. + global	Joint unique monotone	+ Well-nested + resolution + gluing	V_1-V_4
Collapse	Dim. 8 (estab- lished)	Noncomm. (prospective)	Cycle (for- mal)	Scale (no- stic)
\mathcal{O}'	Rectifiable currents	Petz maps + QIG	Quotient cate- gory	Embedd. world-mo agent
Obstruc- tion	Simons cone	Petz violation (cand.)	Minimal cycle	Deploy- ment fa (cand.)
Control param.	Dim. n	Noncomm. η	Load L_c	Model sc
Regime theorem	Reg. (dim \leq 7)	Asympt. fre- quentism (conj.)	Monotone Pop	Token ciency
Status	Completed	Prospective	Formal model	Diagnost

9.3 Open Problems

The quantum Bayes obstruction. Does Conjecture 6.5 hold? A simulation campaign varying Hilbert-space dimension, channel rank, and commutator norms, testing for sharp threshold behavior in the monotonicity violation rate, would directly address this. A positive result would establish quantum Bayes as a second completed instance of the six-stage pattern.

Spherepop obstruction classification. Classify all minimal obstructions for the standard Spherepop invariant system. Determine whether minimal obstructions have a forbidden-subgraph characterization. Develop regularity theory for the quotient category: under what conditions does the collapsed residue retain sufficient structure for subsequent operations?

New domains. Apply the diagnostic to: the passage from spaces to ∞ -groupoids in homotopy theory; from effective field theories to UV completions in physics; from classical logic to realizability toposes in foundations. In each case, identify the invariant system, the governing operations, and the obstruction structure.

Quantitative accessibility theory. Develop a formal theory measuring the “distance” from the working regime to the obstruction region. Explain why the Simons cone was accessible within a decade while the Petz obstruction (if it exists) has not been formalized in decades. The answer presumably involves the difference between mathematical constructibility (GMT) and computational explorability (quantum Bayes).

Chapter 10

Conclusions

10.1 The Central Claim, Evaluated

The GMT case: all six stages are documented. The invariant system was unsatisfiable in dimension eight. GMT was available. The Simons cone was accessible and provided the witness. The field committed. Regularity theory characterized the recoverable regime. The pattern holds in full.

The quantum Bayes case: Stages 1 through 4 are plausible. Stage 5 is absent. Stage 6 has not occurred. The framework correctly identifies what is missing and specifies what a definitive result would require.

The Spherepop case: the pattern is formally instantiated in a calculus designed to exhibit all six stages explicitly. It provides a formal model for the structural claims, not empirical evidence but a proof of conceptual coherence.

The foundation model case: the framework generates diagnostic questions with specific empirical content rather than resolved answers. This is the appropriate contribution.

10.2 Structure Without Blueprint

The framework developed here is fundamentally anti-teleological. Neither geometric measure theory nor Bayesian inference nor the Spherepop calculus requires a blueprint. Area-minimizing currents do not know the surface they are attempting to become. Bayesian posteriors do not know the parameter they will converge toward. Scope graphs do not contain their future resolutions.

Global structure emerges from local admissibility conditions rather than from pre-specified endpoints. Ontological enlargement is not the discovery of a hidden plan. It is the discovery that the admissibility conditions were richer than previously recognized. The Simons cone does not reveal a deeper blueprint for eight-dimensional geometry. It reveals that the admissibility conditions of area minimization support more structures than the smooth category can contain.

This is the mathematical content of the observation that the universe does not run on blueprints. Stable structure, wherever it appears, is the residue of local admissibility conditions operating without predetermined targets.

10.3 Obstructions Do Not Create Enlargement

Mathematical progress is often described as the discovery of new objects. The framework developed here suggests a

different picture.

New objects matter because they reveal that old ontologies were not closed under their own governing dynamics. The Simons cone did not discover a new kind of geometry. It revealed that smooth geometry was already insufficient for the questions area minimization had always been asking. The Petz map does not discover a new kind of probability. It reveals that classical probability was already insufficient for the inference problems that quantum mechanics poses. A collapse cycle in a scope graph does not create an irreducible dependency. It reveals that the dependency structure was already irresolvable within the current containment hierarchy.

Obstructions do not create enlargement. They make enlargement unavoidable.

This is why a single object can reorganize a field. Not because it is exotic or complicated, but because it certifies, by its existence, that the field was already living in a larger world than it knew.

We do not fade away. We resonate, as stable configurations within dynamics that were always richer than our descriptions of them. The old ontologies do not disappear. They persist as regimes, embedded within the larger structures that contain them, resonating at the frequencies where the governing operations remain within the recoverable subclass.

Return is possible. Not return to the old ontology as a universal. Return to the conditions under which it remains valid. That is what regularity theory is. That is what decoherence

theory is. That is what scope completion is. The enlarged framework is valuable precisely because it tells us when we can go home.

References

- [Alm66] Frederick J. Almgren. “Some interior regularity theorems for minimal surfaces and an extension of Bernstein’s theorem”. In: *Annals of Mathematics* 84 (1966), pp. 277–292.
- [Ber15] Serge Bernstein. “Sur un théorème de géométrie et ses applications aux équations aux dérivées partielles du type elliptique”. In: *Communications de la Société Mathématique de Kharkov* 15 (1915), pp. 38–45.
- [BDG69] Enrico Bombieri, Ennio De Giorgi, and Enrico Giusti. “Minimal cones and the Bernstein problem”. In: *Inventiones Mathematicae* 7 (1969), pp. 243–268.
- [Che82] Nikolai N. Chentsov. “Statistical Decision Rules and Optimal Inference”. In: *Translations of Mathematical Monographs* 53 (1982).
- [De 65] Ennio De Giorgi. “Una estensione del teorema di Bernstein”. In: *Annali della Scuola Normale Superiore di Pisa* 19 (1965), pp. 79–85.
- [Fed69] Herbert Federer. *Geometric Measure Theory*. New York: Springer, 1969.

- [Pet88] Dénes Petz. “Sufficiency of channels over von Neumann algebras”. In: *Quarterly Journal of Mathematics* 39.1 (1988), pp. 97–108.
- [Sch32] Erwin Schrödinger. “Sur la théorie relativiste de lélectron et línterprétation de la mécanique quantique”. In: *Annales de l’Institut Henri Poincaré* 2 (1932), pp. 269–310.
- [Sim83] Leon Simon. *Lectures on Geometric Measure Theory*. Canberra: Australian National University, 1983.
- [Sim68] James Simons. “Minimal varieties in Riemannian manifolds”. In: *Annals of Mathematics* 88.1 (1968), pp. 62–105.
- [Sin67] Richard Sinkhorn. “Diagonal equivalence to matrices with prescribed row and column sums”. In: *American Mathematical Monthly* 74.4 (1967), pp. 402–405.
- [WWY15] Mark M. Wilde, Andreas Winter, and Dong Yang. “Strong converse for the classical capacity of entanglement-breaking and Hadamard channels via a sandwiched Rényi relative entropy”. In: *Communications in Mathematical Physics* 331 (2015), pp. 593–622.