Foundations of Distinction Geometry

A Unified Catalogue of the Flyxion Research Program

Flyxion

Independent Researcher

June 23, 2026

A unified framework showing how distinctions, admissibility, repair, projection, and coordination arise from a single architecture.

"What exists?"

$\longrightarrow$

"What distinctions survive projection?"

"A system can only act, remember, predict, coordinate, or persist to the extent that it maintains distinctions."

Contents

Preface

This document is both a catalogue and an argument. The catalogue records the frameworks, theories, formal systems, and computational artefacts developed under the name Flyxion over roughly a decade of independent research. The argument is that they are not separate inventions. They are cross-sections through a single architecture, and that architecture has a name: distinction geometry.

The central observation is simple. A system can only act, remember, predict, coordinate, or persist to the extent that it maintains distinctions. When distinctions collapse — through noise, compression, forgetting, institutional drift, physical coarse-graining, or the inherent many-to-one character of any observation — something is lost that cannot be recovered without additional information. The geometry of that loss, and the geometry of what survives, is the subject of this program.

The frameworks documented here include: **RSVP** (Relativistic Scalar-Vector Plenum, a field-theoretic cosmology grounded in constraint-first dynamics); **CLIO** (constraint-layered inference and projection theory); **TARTAN** (trajectory-aware recursive tiling with annotated noise); **HYDRA** (coordination geometry and collective admissibility as categorical colimit); **Spherepop** (an irreversible event calculus with Pop, Refuse, Bind, and Collapse operators); the **Admissibility Program** (distinguishability geometry and reachability-based ontology); the **Repair Theory** (repair as a primitive prior to optimization); **Distinction Geometry** (Jensen–Shannon interface metrics and the curvature of distinguishability); the **Preference Field Program**; and the **Ecology of Distinctions** (a 31-chapter academic monograph on distinction theory, reachability, and admissibility geometry).

Earlier works often asked: *What structures exist in the internal space?* The current program asks: *What can be known about the internal space from observations?* And then takes the further step: *Perhaps the geometry of that limitation is itself the primary object.*

This document is organized so that a reader beginning with Chapter 1 arrives at quadratic gravity, coordination geometry, the emergence of quantum probability, and the geometry of scientific revolutions as consequences of the same underlying architecture, rather than as dozens of independent projects.

Flyxion
Independent Researcher

Part I

The Primitive Problem

Chapter 1

Why Distinctions Matter

1.1. The Central Observation

A system can only act, remember, predict, coordinate, or persist to the extent that it maintains distinctions. This is not a claim about information theory, or about physics, or about cognition in particular. It is a claim about what any system must do to be a system at all.

Consider action. To act differently in different circumstances, a system must be able to distinguish those circumstances. A thermostat that cannot distinguish hot from cold cannot regulate temperature. A market that cannot distinguish solvent from insolvent firms cannot allocate capital. A nervous system that cannot distinguish threat from safety cannot protect the organism it inhabits. Action requires discrimination, and discrimination requires maintained distinction.

Consider memory. To remember is to preserve a distinction between what was and what is, between an earlier state and the current one. A system with no distinctions has no memory, not because it lacks storage, but because it cannot mark one time as different from another. The content of memory is the set of distinctions that have been preserved across time.

Consider prediction. To predict is to distinguish possible futures from one another and from the present. A prediction that cannot distinguish outcomes carries no information about what will happen. The quality of a predictive model is precisely the quality of the distinctions it can maintain across time and circumstance.

Consider coordination. For two systems to coordinate, they must share distinctions: they must agree on what counts as the same and what counts as different. Language is a technology for sharing distinctions. Law is a technology for enforcing them. Mathematics is a technology for making them precise.

Consider persistence. A system persists when its identity-constituting distinctions are maintained against the pressures that would collapse them. An organism persists when it maintains the distinction between self and environment. An institution persists when it maintains the distinctions between its roles, its procedures, and its goals. A physical structure persists when the forces that would homogenize it are outweighed by the forces that maintain its differentiation.

Principle 1.1 (Distinction Priority)**.** Distinctions are prior to information, entropy, memory, computation, agency, knowledge, and coordination. Each of these concepts can be derived from the notion of maintained distinction; none of them can be defined without it.

1.2. What Distinction Collapse Looks Like

The collapse of distinctions is one of the most pervasive phenomena in nature and in human affairs, and it takes many forms.

Physical collapse. Thermodynamic equilibrium is the state in which all thermodynamic distinctions within a system have been eliminated: temperature, pressure, and chemical potential are uniform everywhere. The approach to equilibrium is the approach to distinction collapse. Entropy, in this reading, is a measure of how many distinctions have been lost: how many microstates are compatible with the current macrostate, how large the preimage of the current observation.

Cognitive collapse. Forgetting is the collapse of temporal distinctions: the inability to mark what was as different from what is. Cognitive biases are often distinction collapses: the availability heuristic collapses the distinction between frequency and memorability; the halo effect collapses the distinction between one trait and all others; confirmation bias collapses the distinction between evidence and expectation.

Institutional collapse. Goodhart's Law is a distinction collapse: when a measure becomes a target, the distinction between the measure and the thing it was measuring collapses. Goal drift is a distinction collapse: the distinction between the stated purpose of an institution and its actual behavior erodes over time. Bureaucratic ossification is a distinction collapse: the distinction between the form and the substance of a procedure disappears, leaving only the form.

Scientific collapse. Paradigm exhaustion, in Kuhn's sense, is a distinction collapse: the reigning framework can no longer distinguish anomalies from noise, can no longer separate what it predicts from what it observes. A scientific revolution is a distinction reorganization: new categories are introduced that make previously invisible differences visible.

Computational collapse. Information loss in computation is a distinction collapse. Irreversible operations — AND, OR, NAND — map multiple input states to a single output state, collapsing the distinction between inputs. The thermodynamic cost of irreversible computation is the energy required to erase the record of that collapsed distinction.

1.3. Why This is the Right Starting Point

The history of the foundational sciences is a history of discovering that something thought to be primitive is in fact derived from something more basic. Newtonian mechanics was derived from the calculus of variations. Thermodynamics was derived from statistical mechanics. Information theory was derived from probability. Probability is currently being rederived from admissibility and projection.

The Flyxion program proposes that distinctions are the next primitive: that information, entropy, probability, and admissibility are all derived from the more basic notion of maintained distinction. This is not a speculative claim. The derivations are largely in hand; they are documented in the subsequent chapters of this work.

What is primitive is not an object, a state, a field, or a wave function. What is primitive is the capacity to tell things apart.

Chapter 2

Distinction Spaces

2.1. Formal Definition

Definition 2.1 (Distinction Space)**.** *A* distinction space *is a pair* $(X, \sim)$ *where* X *is a set and* $\sim$ *is an equivalence relation on* X*. An element* $x \in X$ *is a* history*; the equivalence class* $[x]_\sim$ *is the* observable class *of* x*. Two histories* $x, x' \in X$ *are* indistinguishable *if* $x \sim x'$ *and* distinguishable *otherwise.*

Every equivalence relation on X induces a projection

$$\Pi : X \to X/\sim \tag{2.1}$$

onto the quotient set of equivalence classes. The projection Π records everything that can be observed; the fiber $\Pi^{-1}([x])$ records everything that cannot.

Remark 2.1. This is a very general setup. When X is the microstate space of a thermodynamic system and $\sim$ is thermodynamic coarse-graining, we recover the foundations of statistical mechanics. When X is the space of neural activation patterns and $\sim$ is behavioral equivalence, we recover a framework for cognitive science. When X is the space of gauge field configurations and $\sim$ is gauge equivalence, we recover the fiber bundle structure of Yang–Mills theory. The apparatus of distinction spaces is not a metaphor across these domains; it is the same mathematical object.

2.2. Distinguishability Classes and Their Properties

The quotient $X/\sim$ is the space of *distinguishability classes*: the collection of all things that can be told apart. It inherits whatever structure X carries that is compatible with $\sim$. When X is a topological space and $\sim$ is an open equivalence relation, $X/\sim$ is a topological space. When X is a measure space and $\sim$ is measurable, $X/\sim$ is a measure space.

Definition 2.2 (Refinement and Coarsening)**.** *A distinction space* $(X, \sim')$ *is a* refinement *of* $(X, \sim)$ *if* $x \sim' y \Rightarrow x \sim y$*: finer equivalences are available. It is a* coarsening *if* $x \sim y \Rightarrow x \sim' y$*: fewer distinctions are maintained. Coarsening is projection; refinement is distinction creation.*

Definition 2.3 (Representational Entropy)**.** *Given a distinction space* $(X, \sim, \mu)$ *with reference measure* μ*, the* representational entropy *of the equivalence relation* $\sim$ *is*

$$H(\sim \mid \mu) = -\sum_{[x] \in X/\sim} \mu([x]) \log \mu([x]). \tag{2.2}$$

This measures the uncertainty about which equivalence class a randomly drawn history belongs to. It is maximised when all classes are equally likely (maximum distinction preservation) and minimised when all histories are in a single class (complete distinction collapse).

2.3. Projection, Collapse, and Compression

The three fundamental operations on distinction spaces are projection, collapse, and compression. They are related but distinct.

Projection is the map $\Pi : X \to X/\sim$. It sends each history to its equivalence class. It is the operation of observation: it records what is visible and discards what lies in the fiber.

Collapse is a change in the equivalence relation that coarsens it: $\sim$ is replaced by a coarser $\sim'$. It corresponds to the process by which a system loses the ability to distinguish things it previously could distinguish. Forgetting is cognitive collapse. Equilibration is thermodynamic collapse. Goal drift is institutional collapse.

Compression is a deliberate coarsening that preserves a specified subset of distinctions while collapsing others. Efficient compression preserves the distinctions that matter for a given purpose while eliminating those that do not. The theory of compression is therefore the theory of which distinctions matter for which purposes — a question that connects compression theory to decision theory, value theory, and the theory of admissibility.

Chapter 3

The Lifecycle of Distinctions

3.1. Creation

Distinctions are created when a system acquires the capacity to tell things apart that it previously could not. This happens through:

Differentiation. A homogeneous system spontaneously breaks into distinguishable regions. Phase transitions in physics, cell differentiation in biology, and the emergence of social roles in institutions are all examples of distinction creation through differentiation.

Learning. A cognitive system updates its projection Π to distinguish stimuli that it previously conflated. Perceptual learning — the ability to distinguish wines, musical intervals, or faces with increasing precision — is distinction creation.

Measurement. The construction of a measuring instrument creates distinctions by coupling a physical system to a record-keeping apparatus whose states correspond to the distinctions being introduced. The history of science is substantially the history of constructing instruments that create previously unavailable distinctions.

Language. The introduction of a new word or concept creates a distinction by providing a label that separates what was previously unseparated. Technical vocabulary is distinction-creation technology.

3.2. Preservation

A distinction is preserved when the processes that would collapse it — noise, entropy, drift, forgetting — are outweighed by the processes that maintain it. Preservation requires active work.

The energy cost of maintaining a distinction is a central concern of Landauer's principle: erasing one bit of information requires a minimum of $k_B T \ln 2$ joules. Maintaining a distinction against thermal noise requires a comparable investment. This connects distinction preservation to thermodynamics and to the energetics of life.

The social cost of maintaining a distinction is equally real. Legal distinctions require enforcement. Linguistic distinctions require use. Institutional distinctions require active reproduction through practice. A distinction that is not exercised tends to collapse.

3.3. Repair

Repair is the restoration of a distinction that has been partially or fully collapsed. It is the process of recovering information about the generating history from its observable image, when that image is insufficient to determine the history uniquely.

Principle 3.1 (Repair Priority)**.** Repair is prior to optimization, prediction, and intelligence. A system that cannot repair distinctions that have been collapsed cannot learn from experience, cannot correct errors, and cannot adapt to novelty. Repair is the capacity that makes all other cognitive capacities possible.

Repair takes many forms: error correction in coding theory, memory retrieval in cognitive science, paradigm revision in epistemology, immune response in biology, constitutional amendment in political theory. All of these are instances of the same operation: the partial restoration of a collapsed distinction from whatever information survives.

Repair Theory is developed in Part IV of this work.

3.4. Loss

Distinction loss is irreversible in the precise sense established by Distinction Geometry: when the projection Π is non-injective and the lost information is not recorded elsewhere, no reconstruction map can perfectly recover the original history. The minimum reconstruction defect is positive.

This irreversibility is not a failure of ingenuity. It is a theorem. The theorem is what gives the second law of thermodynamics its force: the coarse-graining projection from microstates to macrostates is non-injective, and the minimum reconstruction defect grows with time under any measure-preserving dynamics.

3.5. Reorganization

Reorganization is the replacement of one distinction space $(X, \sim)$ by another $(X, \sim')$, where $\sim'$ is neither a refinement nor a coarsening of $\sim$: it makes distinctions that $\sim$ did not make and collapses distinctions that $\sim$ preserved.

Scientific revolutions are distinction reorganizations. The Copernican revolution replaced a distinction space in which the earth and the heavens were distinguished by a distinction space in which the earth is a planet like any other. The quantum revolution replaced a distinction space in which position and momentum were simultaneously sharp by one in which they are fundamentally conjugate.

Conceptual change, paradigm shifts, learning, and creativity are all, in the present framework, instances of distinction reorganization. They are the topic of Chapter 21.

Part II

Information and Projection

Chapter 4

Information as Distinction

4.1. Deriving Information Theory

The standard foundations of information theory take probability as primitive and derive entropy and information from it. The distinction geometry program reverses this: distinctions are primitive, and information is derived from them.

Definition 4.1 (Information as Maintained Separation)**.** *The* information content *of a distinction between x and y in a distinction space $(X, \sim, \mu)$ is the negative log-probability of the event that a randomly drawn pair (x', y') belongs to the same equivalence class:*

$$I(x \not\sim y) = -\log P\big([x]_\sim \neq [y]_\sim\big). \tag{4.1}$$

The entropy *of the distinction space is the expected information content: $H(\sim \mid \mu) = \mathbb{E}[I]$.*

Under this derivation, Shannon entropy emerges as the expected distinguishability of a randomly drawn pair: the amount of distinction that the equivalence relation $\sim$ maintains on average. High entropy means many distinctions are maintained; low entropy means few.

Theorem 4.2. *Shannon entropy $H(P) = -\sum_i p_i \log p_i$ equals the expected information content of a uniform-weight distinction space $(X, \sim_P)$ where $\sim_P$ is the equivalence relation induced by the probability distribution P.*

Proof sketch. The probability that a random pair (x, y) drawn i.i.d. from P falls in the same class is $\sum_i p_i^2$. The expected information content of the distinction between them is $-\log(1 - \sum_i p_i^2)$ for distinguishable pairs. Under the standard information-theoretic derivation from the axioms of entropy, the result is $H(P)$ as the unique measure satisfying continuity, symmetry, and the chain rule for conditional distinctions. □

4.2. Compression as Distinction Reduction

Lossless compression is the identification of a coarser representation that preserves all distinctions relevant to a specified purpose. Lossy compression deliberately collapses some distinctions to achieve a more compact representation. The rate–distortion tradeoff is the tradeoff between the number of distinctions preserved and the cost of preserving them.

The mutual information $I(X;Y)$ between two random variables is the average amount of distinction about X maintained by Y: the reduction in uncertainty about X given Y. In distinction geometry terms, it is the expected decrease in fiber size induced by conditioning on Y.

4.3. The Data Processing Inequality as a Distinction Theorem

Theorem 4.3 (Data Processing as Distinction Monotonicity)**.** *If $X \to Y \to Z$ is a Markov chain (each subsequent variable is a function of the previous one), then $I(X;Z) \leq I(X;Y)$. Processing can only reduce the distinctions maintained about X.*

The data processing inequality is, in distinction geometry terms, the statement that projection is monotone: every additional projection between the history and the observation can only collapse more distinctions. It is impossible to create distinctions about X by processing Y.

Chapter 5

Projection Geometry: The CLIO Framework

5.1. Overview

The CLIO framework (Constraint-Layered Inference and Projection) provides the formal architecture for projection geometry. It was developed to address a gap in the existing frameworks for inference and learning: the failure to distinguish between what a system can infer from observations (the observable manifold M_{obs}) and what structures the system projects onto the world (the internal manifold M_{int}).

CLIO identifies three layers in any cognitive or inferential system:

(i) The *constraint layer*: the set of admissibility conditions that a representation must satisfy to count as legitimate.

(ii) The *projection layer*: the map $\pi : M_{\text{int}} \to M_{\text{obs}}$ that connects internal representations to observable quantities.

(iii) The *inference layer*: the process by which the system updates its internal representation given new observations, i.e., attempts to invert π.

5.2. The CLIO Projection

The central object of CLIO is the projection

$$\pi : M_{\text{int}} \to M_{\text{obs}}, \tag{5.1}$$

which maps each internal world-model to the set of observations it predicts. Two world-models are *projection-equivalent* if they predict the same observations: $m \sim m' \iff \pi(m) = \pi(m')$.

Principle 5.1 (CLIO Projection Sufficiency)**.** No condition on the internal model $m \in M_{\text{int}}$ can be physically or epistemically necessary unless it changes the projection $\pi(m)$. Constraints belong on M_{obs}, not on M_{int}.

This principle is the epistemological foundation for the critique of latent fundamentalism (Chapter 26) and the Observational–Interventional Separation theorem.

5.3. Projection Failure and World-Model Mismatch

A *projection failure* occurs when the internal model m and the observation o are inconsistent: $o \notin \pi(m)$ or $\pi(m) \notin \mathcal{A}_O$. Projection failures come in two kinds:

Structural failure: the internal model's projection misses the actual observation entirely. The model cannot explain what happened.

Admissibility failure: the internal model's projection is consistent with the observation but violates an admissibility constraint. The model explains what happened but in a way that is not admissible (e.g., violates conservation laws, causality, or computational constraints).

Repair Theory (Part IV) is concerned primarily with structural failure: how to update a world-model whose projection no longer includes the current observation.

5.4. Abstraction as Coarsening

Abstraction, in CLIO, is the operation of coarsening the internal model: replacing a fine-grained M_{int} with a coarser M'_{int} in which formerly distinct models are identified. Good abstraction preserves the distinctions that matter for prediction while collapsing those that do not.

The theory of good abstraction is therefore the theory of which distinctions in M_{int} are tracked by distinctions in M_{obs}: which internal distinctions have observable consequences. This is exactly the question answered by the CLIO projection π: the distinctions that matter are those not collapsed by π.

5.5. Latent Fundamentalism and Its Critique

Latent fundamentalism is the error of granting ontological status to internal coordinates merely because they appear in a successful model. It arises when the projection π is forgotten and the internal model is treated as if it were the observable world.

The error is pervasive:

- In quantum mechanics: treating the wave function as a physical object rather than a tool for generating predictions about observations.

- In neuroscience: treating neural activation patterns as the content of experience rather than as projections of that content onto a measurable substrate.
- In economics: treating utility functions as real preferences rather than as models of revealed preferences under projection.
- In institutional theory: treating the organizational chart as the institution rather than as a projection of actual power and decision-making structures.

In each case, the correction is the same: identify the projection π, restrict admissibility conditions to M_{obs}, and treat M_{int} as a representational instrument rather than an ontological object.

Chapter 6

Interface Theory

6.1. The Observable–Internal Distinction

Interface theory is the formal development of the distinction between what is observable and what is internal. Its central claim is that the observable interface of a system — the distribution $P = \pi_*\mu$ on M_{obs} — is the only thing about the system that is empirically accessible.

Definition 6.1 (Observable Interface)**.** *Given a projection setting* (M, O, Π, μ)*, the* observable interface *of the system is the pushforward distribution* $P = \Pi_*\mu \in \mathcal{M}$*, where* $\mathcal{M}$ *is the space of probability distributions on* O*.*

Two systems with the same observable interface are *interface-equivalent*: they cannot be distinguished by any observation, however fine-grained.

6.2. Observational–Interventional Separation

The Observational–Interventional Separation theorem establishes that the observable interface underdetermines the internal structure:

Theorem 6.2 (Observational–Interventional Separation)**.** *The map from internal structures to observable interfaces is not injective: multiple distinct internal structures* $(M^{(1)}, \Pi_1, \mu_1)$ *and* $(M^{(2)}, \Pi_2, \mu_2)$ *may share the same observable interface* $P^{(1)} = P^{(2)}$*.*

This theorem has immediate practical consequences. It establishes that no experiment — however carefully designed — can uniquely determine the internal structure of a system from its observable interface. The choice of internal model is therefore always underdetermined by evidence.

The principle of ontological restraint follows: admissibility conditions should be imposed on observable interfaces, not on internal structures. Internal structures are representational tools; they should be evaluated by their predictive accuracy and computational utility, not by their ontological plausibility.

6.3. Projection Sufficiency Principle

Principle 6.1 (Projection Sufficiency)**.** Physical, cognitive, and institutional predictions depend only on the observable interface $P = \Pi_* \mu$. No condition on the internal structure (M, Π, μ) is necessary unless it changes the observable interface.

This principle unifies the ghost problem in quantum gravity, the underdetermination problem in philosophy of science, the model selection problem in statistics, and the legibility problem in institutional theory. In each case, the operative principle is the same: constraints belong on the interface, not on the internal structure.

6.4. Interface-Native Metrics: The Jensen–Shannon Distance

The appropriate metric for comparing observable interfaces is the Jensen–Shannon distance, developed in detail in Distinction Geometry (Chapter 27). Here we record its key properties.

Given two observable interfaces $P, Q \in \mathcal{M}$, the distinction functional is

$$\mathcal{D}(P, Q) = H\big(\tfrac{P+Q}{2}\big) - \tfrac{1}{2}\big(H(P) + H(Q)\big) = \mathrm{JSD}(P\|Q). \tag{6.1}$$

The interface distance is $d_{\mathrm{int}}(P, Q) = \mathcal{D}^{1/2}(P, Q)$.

It is interface-native: its definition requires no reference to any internal structure. It is symmetric, satisfies the triangle inequality, and vanishes if and only if $P = Q$. It is the unique metric on $\mathcal{M}$ that is interface-native, symmetric, and locally equivalent to the Fisher information metric.

6.5. The Derived Interface Conjecture

The Derived Interface Conjecture is the central unifying claim of the physical applications of the framework:

Conjecture 6.1 (Derived Interface)**.** Observable spacetime theories are effective interfaces generated by projection from a larger admissibility manifold. Gauge symmetry, Hilbert structure, and metric positivity are properties of the interface rather than of the underlying admissibility structure. Two theories are physically equivalent if and only if $d_{\mathrm{int}} = 0$ across all observable energy scales.

In particular, Hilbert-space quantum gravity and Krein-space quadratic gravity are conjectured to satisfy $d_{\mathrm{int}} = 0$ below the Planck scale.

Part III

Admissibility

Chapter 7

The Admissibility Program

7.1. Core Idea

The Admissibility Program asks: which histories, configurations, or trajectories of a system are physically, cognitively, or institutionally *realizable*? The answer is not given by a list of allowed states but by a set of constraint closure conditions: a configuration is admissible if it satisfies a closed system of mutual constraints.

Definition 7.1 (Admissibility Manifold)**.** *An* admissibility manifold *is a triple* $(M, \mathcal{A}, \Pi)$ *where* M *is a space of histories,* $\mathcal{A} \subseteq M$ *is the admissible subset (those satisfying all constraint closure conditions), and* $\Pi : M \to O$ *is the projection to the observable interface.*

The admissible subset $\mathcal{A}$ is determined by the constraint structure of the theory: which histories satisfy all the equations of motion, boundary conditions, conservation laws, and causal ordering conditions simultaneously.

7.2. Admissibility Geometry

The admissibility manifold has a natural geometry: the metric on M that measures how far a history is from the admissibility boundary $\partial\mathcal{A}$.

Definition 7.2 (Admissibility Distance)**.** *The* admissibility distance *of a history* $h \in M$ *is*

$$d_{\mathcal{A}}(h) = \inf_{h' \in \mathcal{A}} d(h, h'), \tag{7.1}$$

where d *is a metric on* M*. A history is admissible iff* $d_{\mathcal{A}}(h) = 0$*.*

Definition 7.3 (Admissibility Distortion)**.** *The* admissibility distortion *of a theory perturbation* δT *is the change in the admissibility manifold* $\mathcal{A}$ *induced by* δT*: how the boundary of admissibility shifts.*

7.3. Constraint Closure

A system is *constraint-closed* if its admissibility conditions are mutually consistent: satisfying one constraint does not force violation of another. Constraint closure is the condition that the admissibility manifold $\mathcal{A}$ is non-empty.

Constraint closure failure is the signature of a contradictory or over-determined theory: a theory that requires more of its histories than any history can provide. The history of physics contains several examples: the original formulation of Dirac's equation required both positive-energy and negative-energy solutions (constraint closure was restored by the hole theory and later by quantum field theory); the original formulation of the Einstein–Hilbert action was non-renormalizable (constraint closure at the quantum level required additional terms).

7.4. Admissibility Equivalence and Theory Space

Two admissibility manifolds $(M^{(1)}, \mathcal{A}^{(1)}, \Pi_1)$ and $(M^{(2)}, \mathcal{A}^{(2)}, \Pi_2)$ are *admissibility-equivalent* if $\Pi_1(\mathcal{A}^{(1)}) = \Pi_2(\mathcal{A}^{(2)})$: their projections to the observable interface agree on all admissible histories.

This is a weaker condition than structural equivalence (isomorphism of the internal manifolds) and a stronger condition than empirical equivalence (agreement on all observed outcomes). Admissibility equivalence is the correct notion of theoretical equivalence for the purposes of the present program.

Chapter 8

Reachability and Continuation

8.1. Reachability Volumes

The *reachability volume* of a history $h \in \mathcal{A}$ at time t is the volume of the set of admissible future histories consistent with h up to time t:

$$\mathrm{Vol}_t(h) = \mu\big(\{h' \in \mathcal{A} : h'|_{[0,t]} = h|_{[0,t]}\}\big). \tag{8.1}$$

A history with large reachability volume has many admissible continuations; a history with small reachability volume is nearly determined by its past.

Reachability volume is the admissibility-theoretic analogue of entropy: it measures how many futures are open from a given past. The second law of thermodynamics, in this framework, becomes the claim that reachability volume is non-decreasing under the macroscopic dynamics, which follows from the non-injectivity of the coarse-graining projection.

8.2. Boundary Effects and the Xylomorphic Criterion

The boundary of the admissibility manifold $\partial\mathcal{A}$ plays a special role: histories near the boundary have few admissible continuations. A system near $\partial\mathcal{A}$ is fragile: small perturbations may push it outside the admissible region entirely.

The *xylomorphic criterion* is the condition $\lambda < 1$ for regenerative systems: a system is regenerative (capable of self-repair and continuation) if its deformation rate is less than its regeneration rate. Systems satisfying the xylomorphic criterion maintain a positive distance from $\partial\mathcal{A}$; systems violating it drift toward the boundary and eventual collapse.

8.3. Path Dependence

The admissibility manifold may be path-dependent: whether a history h is admissible may depend not only on its current state but on how it arrived there. Path dependence introduces irreversibility into the admissibility structure itself, not just into the dynamics.

Path dependence is the formal correlate of hysteresis in physics, institutional lock-in in social theory, and traumatic memory in cognitive science. In each case, the past constrains the admissible futures in ways that cannot be fully captured by the current state alone.

8.4. Continuation versus Termination

A *continuation* is a history that extends an existing admissible history while remaining admissible. A *termination* is a history that has no admissible continuation: it has reached a boundary of $\mathcal{A}$ from which there is no return.

The question of whether a given history has continuations is the admissibility-theoretic version of the survival question: can this system, process, or institution continue? The theory of continuation connects the admissibility program to the biology of persistence, the thermodynamics of dissipative structures, and the political theory of state legitimacy.

Chapter 9

Preference Fields

9.1. Preference Geometry

The Preference Field Program formalizes the geometry of value, preference, and decision. Its central claim is that preferences are not a list of desired outcomes but a field defined over the space of possible histories: a function that assigns a preference intensity to each direction in history space.

Definition 9.1 (Preference Field). *A* preference field *is a smooth function $V : \mathcal{A} \to \mathbb{R}$ on the admissibility manifold, together with a gradient flow $\nabla V : \mathcal{A} \to T\mathcal{A}$ pointing in the direction of increasing preference.*

Preference fields generalize utility functions by allowing preferences to depend on the history of the system rather than only on its current state. This makes the framework applicable to systems with path-dependent values: institutions whose goals evolve, cognitive systems whose preferences are formed by experience, and physical systems whose energy landscape shifts with their history.

9.2. Value Landscapes and Decision Surfaces

The level sets of the preference field V are *indifference surfaces*: sets of histories that the system values equally. The gradient of V defines a *decision surface*: the direction in history space that the system will preferentially explore.

A *local optimum* in the preference field is a history $h^* \in \mathcal{A}$ such that $\nabla V(h^*) = 0$ and $\nabla^2 V(h^*)$ is negative definite. It is a history from which any deviation decreases preference. The distinction between local and global optima is central to the theory of learning and adaptation: a system that follows local gradient ascent may become trapped at a local optimum far from the globally preferred history.

9.3. Preference Transport

As the admissibility manifold changes — through learning, experience, or environmental change — the preference field must be transported across the changing landscape. This is the problem of *preference transport*: how do preferences defined relative to one admissibility structure translate to a new one?

Preference transport is the formal correlate of value alignment in artificial intelligence, value change in moral philosophy, and preference revision in decision theory. The tools of parallel transport on manifolds provide a rigorous framework for this problem: the preference field is transported along a path in the space of admissibility manifolds, and the result is a preference field on the new manifold that is as close as possible to the original in the sense of minimizing the total transported difference.

Part IV

Repair Theory

Chapter 10

Repair as a Primitive

10.1. The Priority of Repair

Repair is prior to optimization. A system that cannot repair distinctions that have been collapsed cannot learn from experience, cannot correct errors, and cannot adapt to novelty. Optimization — the improvement of performance along a fixed dimension — presupposes that the system's basic distinctions are intact. Repair is the process that keeps them intact.

Repair is prior to prediction. A prediction requires a stable distinction between what was predicted and what occurred. When that distinction collapses — when the system can no longer tell the difference between its predictions and its observations — prediction becomes impossible. Repair restores the distinction.

Repair is prior to intelligence. Intelligence, in the most general sense, is the capacity to maintain and extend appropriate distinctions in the face of environmental pressure. A system that can only optimize cannot be intelligent in this sense; it can only execute a fixed strategy. A system that can repair can respond to the unexpected.

10.2. Repair versus Optimization

The distinction between repair and optimization is not a matter of degree but of kind. Optimization improves performance along a fixed objective function within a fixed admissibility structure. Repair restores an admissibility structure that has been damaged or a distinction space that has been partially collapsed.

Principle 10.1 (Repair-Optimization Distinction)**.** Optimization operates within a fixed admissibility manifold $\mathcal{A}$ and moves toward the maximum of a preference field V. Repair operates on the admissibility manifold itself: it restores $\mathcal{A}$ when it has been damaged, or restores the distinction space $(X, \sim)$ when distinctions have been collapsed. Repair changes the structure within which optimization operates.

10.3. Formal Definition of Repair

Definition 10.1 (Repair Map)**.** *Given a projection setting* (M, O, Π) *and a distinguished admissibility predicate* $\mathcal{A} \subseteq M$*, a* repair map *for a history* $h \in M \setminus \mathcal{A}$ *is a map* $\rho : M \setminus \mathcal{A} \to \mathcal{A}$ *that returns the closest admissible history:*

$$\rho(h) = \mathrm{argmin}_{h' \in \mathcal{A}} \, d(h, h'). \tag{10.1}$$

The repair cost *of* h *is* $d_{\mathcal{A}}(h) = d(h, \rho(h))$*.*

Repair is not in general unique: there may be multiple closest admissible histories. When repair is non-unique, the system faces a choice about how to restore admissibility, and that choice reflects the system's values and preferences.

10.4. Repair and the Ecology of Distinctions

The ecology of a distinction space is the network of relationships between distinctions: how the preservation or collapse of one distinction affects others. The *Ecology of Distinctions* (a 31-chapter monograph) develops this framework in detail.

Key concepts include: distinction *dependencies* (when maintaining distinction A requires maintaining distinction B); distinction *conflicts* (when maintaining A requires collapsing C); distinction *clusters* (groups of mutually reinforcing distinctions); and distinction *cascades* (when the collapse of one distinction triggers the collapse of many others).

The ecology of distinctions provides the framework for understanding how complex systems — brains, institutions, languages, ecosystems — maintain their coherence against the constant pressure of distinction collapse.

Chapter 11

Memory as Repair

11.1. Memory as Distinction Preservation

Memory is the technology by which a system preserves distinctions through time. A memory is a record that allows the system to mark what was as different from what is, and to use that marking in its current operation.

Definition 11.1 (Memory)**.** *In the projection framework, a* memory *is a state in the internal space M that encodes information about past observations: a history h_{mem} that is observable-equivalent to the past state h_{past} but temporally distinct from it. The* memory fidelity *is $d(h_{past}, \rho(h_{mem}))$, where ρ is the repair map from memory to past state.*

Memory loss is the growth of memory infidelity over time: the increasing reconstruction defect as the memory trace diverges from the original state it encodes.

11.2. The MEM|8 Framework

MEM|8 (Wave-Stabilized Memory Architecture) is a formal framework for persistent memory in systems with dynamic substrates. Its core insight is that memory is stabilized not by the permanence of the substrate but by wave interference: memories are encoded in stable interference patterns in a dynamic medium, analogous to holograms.

Key components of MEM|8:

(i) *Event logs*: explicit records of distinction-preserving transitions in the history of the system.

(ii) *Ecphory*: the process of retrieving a memory by providing partial information (a cue) and completing the distinction pattern.

(iii) *Persistent histories*: the representation of memory as a history in the internal space M rather than as a state.

(iv) *Memory reconstruction*: the inference of the generating history from partial memory traces, using repair maps to fill in collapsed distinctions.

MEM|8 is authored by C. and A. Chenoweth.

11.3. Ecphory and Distinction Completion

Ecphory is the memory process of retrieving a full distinction pattern from a partial cue. In distinction geometry terms, it is a repair operation: the cue provides a partial specification of the target distinction pattern, and ecphory completes the pattern to the nearest admissible full specification.

The quality of ecphory depends on the structure of the distinction space: if the target pattern is in a dense region of $\mathcal{A}$ (many nearby admissible patterns), ecphory will find it readily; if it is in a sparse region, the repair map will produce a nearby but not identical result.

Chapter 12

Intelligence as Distinction Repair

12.1. The Distinction-Repair Account of Intelligence

Intelligence, in the most general sense, is the capacity to maintain and extend appropriate distinctions in the face of environmental pressure that would collapse them. This account subsumes:

Understanding as the capacity to maintain the distinctions relevant to a domain: to keep separate what should be kept separate and to identify as the same what should be identified as the same.

Learning as the capacity to update the distinction space in response to experience: to repair collapsed distinctions, to refine coarse distinctions, and to create new distinctions where none existed.

Scientific discovery as the distinction reorganization that makes previously invisible differences visible: the introduction of a framework in which things that seemed the same are distinguished, and things that seemed different are unified.

Error correction as the repair of collapsed distinctions in a formal system: the restoration of the distinction between correct and incorrect, true and false, valid and invalid.

Anomaly integration as the process of incorporating observations that violate current distinctions into a revised distinction space: the repair of the distinction between observation and expectation.

12.2. Learning as Projection Revision

Learning is, in the CLIO framework, the revision of the projection $\pi : M_{\text{int}} \to M_{\text{obs}}$. A learning event occurs when the current projection fails — when the current internal model predicts an observation that does not occur, or fails to predict one that does. Learning is the update of π (or of the internal model M_{int}) to reduce this mismatch.

Gradient descent is one mechanism for this update: adjust π in the direction that reduces the discrepancy between $\pi(m)$ and the observed o. But gradient descent is not the only mechanism. Paradigm shifts, creative leaps, and abductive inference are also learning mechanisms, and they typically involve a more radical revision of the distinction space than gradient descent can achieve: they are distinction reorganizations rather than distinction refinements.

Part V

Coordination and Society

Chapter 13

Coordination Geometry: The HYDRA Framework

13.1. Overview

The HYDRA framework (Coordination Geometry and Collective Admissibility) formalizes the geometry of coordination: the conditions under which multiple agents with potentially distinct internal models and distinction spaces can act coherently.

The central insight of HYDRA is that coordination is not a property of individual agents but of the *colimit* of their admissibility manifolds: the minimal admissibility structure that contains all individual structures and is consistent with their interactions.

Definition 13.1 (Collective Admissibility). *Given a collection of agents* $(M_i, \mathcal{A}_i, \Pi_i)_{i\in I}$ *and interaction maps* $\phi_{ij} : M_i \to M_j$*, the* collective admissibility manifold *is the categorical colimit*

$$\mathcal{A}_{\text{coll}} = \operatorname{colim}_{i\in I} \mathcal{A}_i, \tag{13.1}$$

computed in the category of measurable spaces with the interaction maps as morphisms. An action profile $(h_i)_{i\in I}$ *is* collectively admissible *if the induced element in* $\mathcal{A}_{\text{coll}}$ *is admissible.*

13.2. Synchronization and the Kuramoto Model

The Kuramoto model of coupled oscillators provides the simplest example of collective admissibility geometry. Each oscillator i has a phase $\theta_i \in [0, 2\pi)$ and a natural frequency ω_i. The interaction among oscillators is given by the coupling

$$\dot{\theta}_i = \omega_i + \frac{K}{N}\sum_{j=1}^{N} \sin(\theta_j - \theta_i), \tag{13.2}$$

where K is the coupling strength. The collectively admissible states are the synchronized states in which all oscillators have the same effective frequency.

In HYDRA terms, each oscillator has an admissibility manifold $\mathcal{A}_i = [0, 2\pi)$ and the collective admissibility manifold is the subspace of phase profiles consistent with synchronization. The transition from incoherence to synchronization at the critical coupling K_c is a phase transition in the collective admissibility geometry.

13.3. Social Coherence and Institutional Alignment

The HYDRA framework extends the Kuramoto intuition to institutions. An institution is a collection of agents with individual admissibility manifolds (individual roles, constraints, and objectives) that interact through interaction maps (reporting structures, information flows, incentive systems).

Institutional coherence is the analog of synchronization: the condition in which the individual admissibility manifolds are aligned enough that collectively admissible action profiles exist. Institutional dysfunction is incoherence: the collective admissibility manifold is empty or very sparse, and coordinated action becomes impossible.

Chapter 14

Institutions as Distinction Maintenance

14.1. Institutions as Distinction-Preservation Technologies

An institution is, at its most basic, a technology for maintaining distinctions that would otherwise collapse. Legal systems maintain the distinction between permitted and forbidden. Market systems maintain the distinction between prices and quantities. Scientific institutions maintain the distinction between evidence and conjecture. Democratic institutions maintain the distinction between legitimate and illegitimate power.

The health of an institution is therefore measurable as its capacity to maintain the distinctions it was designed to preserve. Institutional decay is distinction collapse. Institutional corruption is distinction confusion: the inability to maintain the distinction between the institution's purpose and the interests of those who operate it.

14.2. Legibility and the Hilbert Assumption in Institutions

Scott's notion of *legibility* — the process by which states simplify and standardize the societies they govern — is the institutional form of the Hilbert Assumption: the identification of the institution's representation of a community with the community itself.

When a state maps a complex, evolving, locally structured community onto a legible grid — standard plots, surnames, tax categories, occupational classifications — it creates a simplified model whose projection onto observables (tax records, census data, administrative outcomes) is tractable. But the simplification collapses distinctions that matter to the people in the community, and the resulting policies, calibrated to the simplified model rather than to the community, produce characteristic pathologies.

The framework identifies this as a projection failure: the state's model $\pi(m)$ mismatches the community's observable behavior o, not because the community is irrational but because the state's projection π does not track the relevant distinctions.

14.3. Goodhart Effects as Hidden Admissibility Directions

Goodhart's Law — when a measure becomes a target, it ceases to be a good measure — is, in the framework, an instance of the hidden admissibility direction phenomenon. When an institution optimizes a proxy measure M for a true objective T, it explores a fiber of the projection π : behavior space $\to$ measure space that keeps M fixed while changing T. If the fiber is non-trivial — if behaviors exist that achieve high M while degrading T — optimization along M will find those behaviors.

The remedy is to tighten the projection: to redesign the measurement system so that the fiber of $\pi^{-1}(\text{high } M)$ contains only behaviors that also achieve high T. This is the institutional analog of removing hidden admissibility directions: making the relevant behavior distinctions visible in the measurement system.

Chapter 15

Civilization as a Repair System

15.1. Large-Scale Distinction Preservation

Civilization, at the largest scale, is a system for preserving distinctions against the constant pressure of collapse. The distinctions civilizations preserve include: the distinction between written and lost knowledge (libraries, education systems, scholarly traditions); the distinction between living and dead languages; the distinction between enforced and unenforced law; the distinction between calibrated and uncalibrated measurement; and the distinction between functional and dysfunctional infrastructure.

Each of these is maintained by an active technology: not just a record but a practice of maintaining the record, using it, updating it, and repairing it when it is damaged. Civilization is not an archive. It is a repair system for the distinctions that make further repair possible.

15.2. Science as Distinction Accumulation

Science accumulates distinctions: it creates and preserves the distinctions that allow the world to be described with increasing precision. The distinction between mass and weight; between heat and temperature; between correlation and causation; between statistical significance and practical significance — each of these is a hard-won distinction, created by scientific work and maintained by scientific practice.

Scientific institutions are distinction-maintenance technologies: journals, peer review, replication, and systematic criticism are all mechanisms for preventing the collapse of scientific distinctions into confusion, bias, or fashion.

Part VI

Computation

Chapter 16

Computation as Restriction of Possibility

16.1. The Reachability Account of Computation

Computation, in the reachability framework, is the process of restricting the space of admissible futures. A computation begins with a large admissibility manifold (many possible output states) and proceeds by applying constraint after constraint, each one eliminating inadmissible futures, until only the desired output remains.

Principle 16.1 (Computational Restriction)**.** Computation is not the creation of outputs but the elimination of inadmissible alternatives. A correct computation is one that restricts the admissibility manifold to exactly the correct answer.

This account is the computational analog of the physical insight that constraint-first dynamics is more fundamental than force-first dynamics. In physics, the laws of motion are not forces that push states forward but constraints that select admissible trajectories from all possible ones. In computation, algorithms are not procedures that create outputs but constraint systems that select the admissible output from all possible symbol strings.

16.2. Irreversible Operations and Distinction Collapse

Irreversible logical operations — AND, OR, NAND — collapse distinctions: they map multiple input states to a single output state. The AND gate maps $(0,0)$, $(0,1)$, and $(1,0)$ all to 0, collapsing the distinction between them. The energy cost of this collapse is Landauer's limit: $k_B T \ln 2$ per bit erased.

Reversible computation preserves distinctions: every input state maps to a distinct output state, so the computation is injective. Reversible computation has no

Landauer cost, but it requires more hardware: the intermediate states that irreversible computation discards must be stored in reversible computation.

The distinction geometry framework gives a unified account of both: irreversible computation increases reconstruction defect (the output does not determine the input), while reversible computation preserves it. The tradeoff between energy cost and information retention is the tradeoff between accepting reconstruction defect and paying the energy cost to avoid it.

Chapter 17

Flow Computing: The NCL Framework

17.1. Null Convention Logic

Null Convention Logic (NCL), developed by Karl Fant, is a formalism for asynchronous, self-timed digital circuits. Its key property is that computation is *complete*: the circuit indicates when it has produced a valid output and when it is resetting, eliminating the need for a global clock.

In distinction geometry terms, NCL is a computational formalism in which the circuit maintains the distinction between "computing" and "complete" at all times. The absence of a global clock means that the circuit's operation is not synchronized to an external distinction (the clock cycle) but to an internal one (the completion signal).

NCL circuits are history-first: the output of a gate is not determined by the current input alone but by the history of inputs, in particular by whether the circuit has received a complete, valid input pattern since its last reset. This makes NCL circuits natural candidates for computation on dynamic substrates where the timing of inputs is not controlled.

17.2. Dynamic Substrates

A *dynamic substrate* is a computational medium whose structure changes during computation: its admissibility manifold shifts as the computation proceeds. Neural tissue is a dynamic substrate: synaptic weights change in response to activity, altering the admissibility conditions for future computation. An economy is a dynamic substrate: the rules of exchange, the availability of resources, and the legal framework all change in response to economic activity.

Computation on dynamic substrates requires a formalism in which the distinction space $(X, \sim)$ is not fixed but evolves. The CLIO projection framework handles this by allowing the projection π to change: $\pi_t : M_{\text{int}} \to M_{\text{obs}}$ is a time-dependent projection whose evolution encodes the changing admissibility structure.

Chapter 18

Programs as Histories: The Spherepop Framework

18.1. Overview

Spherepop is an irreversible event calculus in which the fundamental objects are not states but histories. A Spherepop program is a record of events — Pop, Refuse, Bind, Collapse — each of which modifies the admissibility manifold of the computation.

Definition 18.1 (Spherepop Operators)**.** *(i)* ***Pop****: eliminates the distinction between a suspended event and its continuation, resolving a pending computation.*

(ii) ***Refuse****: maintains the distinction between the current state and an inadmissible continuation, blocking a computation that would violate admissibility.*

(iii) ***Bind****: creates a distinction between two events by linking them: one can only occur if the other has occurred.*

(iv) ***Collapse****: irreversibly eliminates all distinctions within a set of alternatives, selecting one and discarding the rest.*

18.2. Histories versus States

The distinction between history-based and state-based computation is central to the Spherepop framework. In state-based computation, the computation is described by a function from states to states: the output is determined by the current state alone. In history-based computation, the output is determined by the entire history of the computation.

History-based computation is more expressive: it can represent dependencies that are invisible in state-based formulations. In particular, Refuse — the operator that blocks inadmissible continuations — cannot be expressed in a purely state-based formalism, because the admissibility of a continuation may depend on the history of how the current state was reached.

Spherepop is implemented as a C interpreter; the formal semantics is developed in the framework of irreversible event calculi.

Part VII

Dynamics and Substrates

Chapter 19

Frozen Processes

19.1. When Moving Variables Become Hidden Constants

A *frozen process* is a dynamic process in which one or more variables that should be treated as evolving have been fixed at a particular value. The fixing may be deliberate (a model simplification) or inadvertent (a failure to notice that the variable changes).

The frozen-process critique is a recurring theme in the Flyxion program. It arises whenever a model replaces a trajectory with a state: whenever what is actually a history is described as a snapshot, and the dynamics of the history is attributed to the snapshot itself.

Example 19.1 (Frozen Time in Physics). The Hamiltonian formalism of classical mechanics "freezes" time: it describes the evolution of a system by a vector field on phase space, treating time as a parameter rather than a dynamical variable. This works well for conservative systems but breaks down when the system's constraints are time-dependent. The frozen-time approximation hides the time-dependence in the constraints, making it invisible to methods that only look at the phase space.

Example 19.2 (Frozen Preferences in Economics). Standard economic models freeze preferences: they treat the utility function as a fixed parameter of the agent. But preferences change in response to experience, advertising, habituation, and social influence. Freezing preferences hides the dynamics of preference formation, making it impossible to model how markets shape the values of their participants.

19.2. Static Models and Their Failure Modes

Static models — models that describe a system by its state rather than its history — fail when the dynamics of the system's history matters for its current behavior. The failure modes are characteristic:

Path dependence appears as irreducible complexity: two systems with the same current state behave differently because they arrived there by different paths. Static models cannot represent this.

Hysteresis appears as inexplicable history-dependence: a system that behaves differently after undergoing a cycle than it did before. Static models can only represent hysteresis by introducing additional state variables, effectively recovering a history-based description.

Emergence appears as behavior that cannot be predicted from the current state of the components: it requires knowing how the components interacted over time. Static models cannot represent true emergence.

Chapter 20

Dynamic Substrates

20.1. Computing on Evolving Manifolds

When the admissibility manifold $\mathcal{A}$ changes over time, both the valid computations and the valid results of those computations change. Computing on a dynamic substrate requires tracking the evolution of $\mathcal{A}$ as well as the state of the computation.

The TARTAN framework (Trajectory-Aware Recursive Tiling with Annotated Noise) addresses this problem. TARTAN represents the admissibility manifold as a tiling of the history space, with each tile annotated by the admissibility conditions that hold within it. As the manifold evolves, the tiling is updated, and the annotations reflect the new conditions.

20.2. Plastic Systems

A system is *plastic* if its admissibility manifold changes in response to its own history: the constraints that govern the system's behavior evolve as the system acts. Neural tissue is plastic: synaptic plasticity changes the admissibility conditions for neural computation in response to activity. Economic systems are plastic: the legal and institutional framework that governs economic activity changes in response to economic outcomes.

Plasticity is the computational correlate of learning: a plastic system updates its own constraint structure in response to experience, thereby changing what is admissible for it in the future. The theory of plasticity is therefore the computational version of the theory of learning, and the distinction geometry of plastic systems is the formal framework within which both can be understood.

Chapter 21

Distinction Reorganization

21.1. Conceptual Change and Paradigm Shifts

A *distinction reorganization* is a change in the distinction space $(X, \sim)$ that is neither a refinement nor a coarsening: it introduces new distinctions and collapses old ones simultaneously. Scientific paradigm shifts are the most dramatic examples, but distinction reorganizations occur at every scale of cognitive and social life.

Example 21.1 (The Copernican Reorganization). The Copernican revolution replaced a distinction space in which the distinction between sublunary and superlunary was fundamental with one in which the distinction between planets and fixed stars replaced it. The old distinction (earthly versus heavenly) was collapsed; the new one (planetary versus stellar motion) was created. This is a genuine distinction reorganization, not a refinement or coarsening.

21.2. Creativity as Distinction Creation

Creativity is, in the framework, the production of new distinctions where none previously existed: the identification of a difference that had not previously been marked, or the introduction of a concept that allows things to be distinguished that were previously conflated.

The geometry of creative acts is the geometry of paths in the space of distinction spaces: a creative act is a movement to a point in the space of equivalence relations on X that was previously unoccupied. The difficulty of creativity is that the space of distinction spaces is vast, and most movements in it are arbitrary. Creative acts are movements to points in this space that are both novel and useful: that create distinctions with high downstream value for repair, coordination, and prediction.

Part VIII

Physical Applications

Chapter 22

RSVP: The Relativistic Scalar-Vector Plenum

22.1. Overview

RSVP (Relativistic Scalar-Vector Plenum) is a field-theoretic cosmology grounded in constraint-first dynamics. Its central claim is that cosmological structure — the distribution of matter, the geometry of spacetime, the arrow of time — emerges from the constraint dynamics of three coupled fields: a scalar field Φ, a vector field $\mathbf{v}$, and an entropy density S.

The RSVP framework is a physical instantiation of the distinction geometry program: the three fields encode, respectively, the intensity, transport, and entropy of distinction-preserving structure in the cosmological medium.

22.2. The Field Equations

The RSVP field equations govern the coupled dynamics of $(\Phi, \mathbf{v}, S)$:

$$\partial_\mu \partial^\mu \Phi + \xi R \Phi = J[\mathbf{v}, S], \tag{22.1}$$

$$\nabla \cdot \mathbf{v} + \partial_t S = \Sigma[\Phi, \mathbf{v}], \tag{22.2}$$

$$\partial_t S + \mathbf{v} \cdot \nabla S = \Pi[\Phi, \mathbf{v}]. \tag{22.3}$$

Here R is the Ricci scalar, J, Σ, and Π are coupling functionals, and ξ is the conformal coupling constant. The scalar Φ drives the geometry through its stress-energy tensor; the vector $\mathbf{v}$ transports distinction-preserving structure; the entropy S measures the local degree of distinction collapse.

22.3. The RSVP Transport Operator and the Born Rule

The RSVP transport operator is

$$Z(x) = \Phi(x)\, e^{\widehat{L}(x)}, \qquad \widehat{L}(x) = \mathbf{v}(x) \cdot \nabla + \theta(x)\widehat{T}, \tag{22.4}$$

where Φ supplies scaling, $\mathbf{v} \cdot \nabla$ supplies transport, and $\theta\widehat{T}$ supplies torsional structure. The observable transition probability between configurations C_i and C_j is

$$T_{ij} = \int_{\gamma: C_i \to C_j} |Z[\gamma]|^2 \, \mathcal{D}\gamma. \tag{22.5}$$

When the coarse-graining projection has fibers admitting a unitary action preserved by RSVP transport, the induced transition matrix is unistochastic: $T_{ij} = |U_{ij}|^2$ for a unitary U. The Born rule of quantum mechanics is recovered as the leading-order approximation when the RSVP dynamics reduces to linear field theory.

22.4. Cosmological Implications

In the cosmological context, RSVP provides a framework for deriving large-scale structure from constraint dynamics rather than from initial conditions. The observed homogeneity, isotropy, and flatness of the universe emerge as consequences of admissibility-flow consistency: the configurations that maximize reachability volume are those of maximum symmetry.

The entropy field S provides a direct connection to the arrow of time: the direction in which S increases is the direction in which distinction-preserving structure is being consumed by the dynamics. This makes the arrow of time an emergent property of the RSVP constraint structure rather than a fundamental asymmetry of the laws of nature.

Chapter 23

Cosmology

23.1. Constraint Geometry in Cosmology

The large-scale geometry of the universe — its homogeneity, isotropy, and spatial flatness — is a constraint geometry problem: which cosmological histories satisfy all the observational and physical constraints simultaneously?

The gravitational entropy program of Boyle and Turok approaches this question from within the projection framework: the observed universe is a typical element of the admissibility manifold of cosmological histories, where typicality is measured by the volume of the admissibility set consistent with the observations. Cosmological simplicity is not a fine-tuned initial condition but a consequence of the admissibility geometry.

23.2. The CPT-Symmetric Universe

The CPT-Symmetric Universe proposal (Boyle, Finn, Turok) provides a projection-based account of the Big Bang. The universe is CPT-symmetric: the post-bang universe is the CPT image of the pre-bang universe, and both are determined by a CPT-invariant vacuum state at the singularity.

In projection terms, the CPT symmetry is a constraint on the admissibility manifold of cosmological histories: only those histories that are CPT-symmetric are admissible. This constraint selects a unique vacuum state, eliminates the need for fine-tuned initial conditions, and determines the density perturbation spectrum from the CPT-symmetric quantum vacuum.

23.3. Entropy and Expansion

The expansion of the universe is, in the RSVP framework, the growth of reachability volume: as the universe expands, the number of admissible future histories grows. The cosmological arrow of time is the direction of reachability growth.

This connects RSVP cosmology to the distinction geometry account of irreversibility: the expansion of the universe is a distinction-creation process, not a distinction-collapse process. The universe creates new distinctions (new regions of spacetime, new physical structures) as it expands; the local thermodynamic arrow of time, which is a distinction-collapse process, is a consequence of local dynamics within an expanding background that is creating distinctions globally.

Chapter 24

Quantum Interfaces

24.1. Projection Geometry in Quantum Mechanics

The quantum mechanical measurement process is a projection from the internal state space (the Hilbert space $\mathcal{H}$ or Krein space $\mathcal{K}$) to the observable space (the spectrum of measurement outcomes). The Born rule is the formula for computing the probability distribution on the observable space induced by this projection.

The central claim of the quantum interface program is that the Born rule is not a fundamental axiom but a consequence of the projection structure. The squaring operation $a \mapsto |a|^2$ is a projection from complex amplitude space to non-negative intensity space; the phase is discarded because it is a hidden direction in the fiber of this projection.

24.2. The Ghost Problem as a Projection Question

The ghost problem in higher-derivative quantum gravity is, in the projection framework, the question of whether negative-norm states in the internal state space M_{adm} are hidden admissibility directions or observable pathologies.

A negative-norm state is a hidden admissibility direction if the projection $\Pi : M_{\text{adm}} \to M_{\text{obs}}$ maps its contribution to zero: if the observable probability distribution is unaffected by the presence of the ghost. The Turok–Bateman proposal — formulated in the Krein-space framework with ghost-parity symmetry — is a claim that the massive ghost graviton of quadratic gravity is a hidden admissibility direction of exactly this kind.

24.3. Krein Spaces and the Observable Interface

A Krein space $\mathcal{K} = \mathcal{K}_+ \oplus \mathcal{K}_-$ with indefinite inner product is a natural M_{adm} for a quantum theory with ghost states. The observable interface is the pushforward distribution on the space of observable outcomes:

$$P(F \mid I) = \mathrm{tr}_+\big((\Pi_F S \Pi_I)^\dagger (\Pi_F S \Pi_I)\big), \tag{24.1}$$

where the trace is taken with respect to the positive Hilbert space structure on $\mathcal{K}_+$. When ghost-parity symmetry holds ($[G, S] = 0$), this probability functional is positive and normalised, even though the underlying space $\mathcal{K}$ has indefinite inner product.

Chapter 25

Derived Interface Physics

25.1. Observable Physics as Projection

The derived interface program is the application of the projection hierarchy to the full structure of physics. Its central claim is that the structural features of physical theories — Hilbert space, gauge symmetry, metric positivity, the Born rule, and spacetime geometry itself — are properties of the observable interface rather than of the underlying admissibility manifold.

Conjecture 25.1 (Derived Interface Physics)**.** Observable spacetime physics is the effective interface generated by projecting an underlying admissibility manifold M_{adm} onto the observable domain M_{obs}. The hierarchy of effective theories — quantum mechanics, quantum field theory, general relativity, quadratic gravity — represents successive approximations to the projection Π at successively higher orders in the admissibility-flow expansion.

25.2. Gauge Symmetry as Fiber Structure

Gauge symmetry is, in the projection framework, fiber structure: the gauge orbit through a field configuration is the fiber $\Pi^{-1}(\Pi(A))$, the set of all configurations that project to the same observable outcome. The gauge symmetry of the theory is the symmetry of the fiber.

This gives gauge symmetry a projective interpretation rather than a fundamental one: gauge symmetry is not a property of the physical world but of the representational choice that introduces more degrees of freedom than the observables require. The physical content of the theory is the projection, not the gauge structure.

25.3. The Admissibility-Flow Expansion

Spacetime curvature, in the derived interface program, is the first-order correction to admissibility flow: a measure of how much the projection Π departs from a globally flat map. The Einstein field equations are the conditions for admissibility-flow consistency at leading order. Quadratic gravity is the second-order term in the expansion. Higher-derivative theories represent higher-order terms, with ghost sectors corresponding to negative-metric directions in the fiber of Π that do not reach the observable interface.

Part IX

Meta-Theory

Chapter 26

Against Latent Fundamentalism

26.1. The Central Error

Latent fundamentalism is the error of granting ontological status to internal coordinates merely because they appear in a successful model. It is the confusion of the instrument with the reality it instruments: the map with the territory, the model with the world, the projection machinery with the observable world it projects onto.

The error is natural and persistent. Every successful model creates a temptation: if this model works, perhaps the objects it refers to are the real objects of the world. The wave function works, so perhaps the wave function is real. The utility function works, so perhaps utility is real. The organizational chart works, so perhaps the organization really is organized that way.

The Flyxion program resists this temptation not by denying that internal structures exist but by insisting on the distinction between what is observable and what is internal. The internal structure is a tool for generating and explaining observable distinctions. It is indispensable. But it is not ontologically primary.

26.2. A Catalogue of Latent Fundamentalist Errors

The following is a partial catalogue of latent fundamentalist errors, each analyzed in terms of the projection framework:

Wave function realism: treating the wave function as a physical object rather than as an element of the internal space M_{adm} from which observable probabilities are projected.

Utility realism: treating the utility function as a real feature of agents rather than as a projection of their revealed preferences onto a mathematical representation.

Organizational chart realism: treating the organizational chart as the institution rather than as a projection of actual decision-making and power structures.

Metric realism in quantum gravity: treating the positive-definiteness of the Hilbert inner product as a necessary feature of any physically admissible theory rather than as a sufficient condition for observable consistency.

GDP fundamentalism: treating GDP as a measure of welfare rather than as a projection of economic activity onto a single number that collapses the distinction between types of activity that matter for welfare in very different ways.

26.3. The Correction

The correction for latent fundamentalism is always the same: identify the projection Π, locate the observable interface M_{obs}, and restrict admissibility conditions to that interface. Internal coordinates are then free to be whatever they need to be to generate the correct observable predictions, without those coordinates being awarded ontological status.

This is not instrumentalism: the claim is not that internal structure is unreal or that we should be agnostic about it. The claim is that internal structure is epistemologically secondary: the appropriate object of inquiry is the geometry of the observable interface, and internal structures are evaluated by their utility for generating and explaining that geometry.

Chapter 27

Distinction Geometry

27.1. The Interface-Native Metric

Distinction Geometry is the formal development of the geometry of distinguishability between observable interfaces. Its primitive is the Jensen–Shannon distance

$$d_{\text{int}}(P, Q) = \text{JSD}^{1/2}(P\|Q) = \left[H\left(\tfrac{P+Q}{2}\right) - \tfrac{1}{2}\left(H(P) + H(Q)\right)\right]^{1/2}, \tag{27.1}$$

which compares observable interfaces without reference to any underlying admissibility manifold.

The distinction functional $\mathcal{D}(P, Q) = \text{JSD}(P\|Q)$ measures the information destroyed when the distinction between P and Q is forgotten: the entropy of the merged interface minus the average entropy of the separate interfaces. This is the information that disappears when an observer stops distinguishing between the two theories.

27.2. Theory Space and Its Curvature

The space of theories $(\mathcal{M}, d_{\text{int}})$ is a metric space that is locally Riemannian with the Fisher information metric. Its curvature has a precise interpretation in terms of the difficulty of theoretical inference:

Positive sectional curvature at a theory P means that nearby theories are more similar in their observable predictions than their parameter-space distance suggests: many distinct internal models collapse onto nearly identical observables. This is the geometry of overfitting, paradigm stability, and institutional opacity.

Negative sectional curvature at a theory P means that small changes in the theory produce large changes in observable predictions: the theory is highly falsifiable in that parameter direction.

Zero curvature is the Euclidean regime: parameter distance and observable distance scale proportionally, and the theory is locally well-identified.

27.3. Reconstruction Defect and Irreversibility

The reconstruction defect $\Delta_R(h) = d(h, R(\Pi(h)))$ measures the minimum error that any reconstruction must make in recovering the generating history from an observation. When the projection Π is non-injective, no reconstruction map achieves zero defect: irreversibility is a theorem about the structure of projection, not a contingent feature of dynamics.

This unifies the arrow of time, thermodynamic irreversibility, memory loss, scientific underdetermination, and institutional goal drift as instances of a single phenomenon: persistent reconstruction defect under non-injective projection.

Chapter 28

The Ecology of Distinctions

28.1. Overview

The Ecology of Distinctions is a 31-chapter academic monograph (current length: approximately 30,000 lines of LaTeX) that develops the full formal theory of how distinctions interact, support each other, compete with each other, and form systems. It is the most detailed formal development of the distinction program.

The monograph covers: formal distinction theory and its connection to equivalence relations; reachability in distinction spaces; admissibility geometry; the boundary theory of distinction collapse; repair algebras for the restoration of collapsed distinctions; distinction ecology (the network structure of distinction dependencies and conflicts); biological applications (cell types as distinction spaces, morphogenesis as distinction creation, homeostasis as distinction preservation); and a full treatment of the admissibility manifold hierarchy including proofs of the Projection Pathology Separation Theorem and the Observable Admissibility Theorem.

28.2. The Master Chain

The Ecology of Distinctions concludes with the following chain, which is the organizing principle of the entire program:

$$\text{Distinction} \to \text{Information} \to \text{Entropy} \to \text{Repair} \to \text{Continuation} \to \text{Admissibility} \to \text{World} \tag{28.1}$$

Distinctions enable information. Information enables entropy accounting. Entropy accounting enables repair. Repair enables continuation. Continuation generates admissibility conditions. Admissibility conditions, accumulated over histories, generate the structured world.

Part X

Unified Framework

Chapter 29

The Master Dependency Chain

29.1. From Distinctions to Worlds

The architectures described in the preceding parts are not independent inventions. Each is a level in a single dependency chain. The chain goes:

Distinctions are the primitive. A system that can tell things apart can do everything else that follows. A system that cannot is indistinguishable, by any observation, from a system with no internal structure at all.

Information is maintained distinction. The Shannon entropy of a source is the expected distinguishability it produces. Mutual information is the distinctions shared between two sources. Compression is distinguishability reduction. Information theory is the theory of distinctions in the presence of uncertainty.

Entropy is collapsed distinction. Thermodynamic entropy measures the number of microstates compatible with a macrostate — the size of the fiber of the coarse-graining projection. High entropy means many internal histories produce the same observable, which means many distinctions have been collapsed by the projection. The second law is the statement that, under the dynamics, the fiber grows.

Repair is distinction restoration. When a distinction has been collapsed — when the fiber has grown, when information has been lost, when the gap between the internal history and its observable image has increased — repair is the attempt to restore the collapsed distinction from whatever information remains. Repair is possible only to the extent that the collapsed distinction left traces in the observable world; it is impossible to the extent that the collapse was complete.

Continuation is the capacity of a system to extend its history into the future while remaining admissible. A system that has lost too many distinctions — whose reachability volume has collapsed too far — cannot continue: it has reached a boundary of $\mathcal{A}$

from which there is no admissible exit. Repair maintains the distinctions that make continuation possible.

Admissibility is the accumulated constraint closure of a history: the set of conditions that a history must satisfy to count as realizable. Admissibility is generated by the interaction of distinctions over time: each distinction imposes constraints on which subsequent distinctions are possible. The admissibility manifold is the set of all histories that satisfy all constraints simultaneously.

Worlds are admissibility manifolds that have achieved the kind of closure, continuity, and coherence that we associate with a persistent, interacting domain of reality. A world, in this sense, is a distinction-preservation system that has become self-sustaining: its distinctions generate the constraints that maintain those distinctions.

29.2. The Role of Each Framework

Within this chain, each framework in the Flyxion program occupies a specific role:

RSVP instantiates the chain in a physical field theory: Φ carries distinction intensity, $\mathbf{v}$ carries distinction transport, and S carries distinction entropy. The RSVP field equations are the physical laws of distinction flow.

CLIO formalizes the projection layer: it is the theory of how internal distinction spaces map onto observable ones, and what can be inferred about the former from the latter.

TARTAN handles computation on dynamic substrates: the case where the admissibility manifold changes during the process of distinction maintenance.

HYDRA extends the framework to collectives: the case where multiple distinction spaces must be coordinated into a collective admissibility manifold.

MEM|8 formalizes the memory layer: the technology by which a system preserves distinctions through time against the pressure of collapse.

Spherepop provides the computational semantics: a formalism for processes that are irreversible, history-dependent, and distinction-collapsing.

The Admissibility Program develops the geometry of admissibility manifolds: the formal theory of which histories are realizable and how the space of realizable histories is structured.

Repair Theory develops the algebra of distinction restoration: the formal theory of how collapsed distinctions are partially restored from surviving traces.

Distinction Geometry provides the interface metric: the formal theory of how observable interfaces are compared, how theories are distinguished from each other, and what the geometry of that comparison looks like.

Coordination Geometry extends the projection framework to multi-agent systems: the formal theory of when agents with different distinction spaces can achieve collective admissibility.

Preference Fields provide the value layer: the formal theory of how preferences over distinctions are represented, transported, and traded off.

Chapter 30

Open Problems

30.1. Mathematical Open Problems

The following mathematical problems arise directly from the frameworks documented in this work and have not yet been fully resolved:

Distinction Geometry

Completeness of the theory space. Is the metric space $(\mathcal{M}, d_{\text{int}})$ complete? The question is whether Cauchy sequences of observable interfaces converge to an interface that can be realised by some admissibility manifold.

Uniqueness of the JSD metric. Is $d_{\text{int}} = \text{JSD}^{1/2}$ the unique interface-native metric satisfying symmetry, triangle inequality, and local Fisher-metric equivalence? An axiomatic characterization would clarify its status as the canonical metric on theory space.

Curvature bounds. What bounds on the sectional curvature of $(\mathcal{M}, g^F)$ are implied by admissibility conditions? In particular, does high admissibility (a small admissibility manifold relative to the full history space) imply high curvature of the theory space?

Repair algebras. What algebraic structures do repair maps form? In particular: is the composition of two repair maps a repair map? Is there an identity repair? What is the appropriate notion of morphism between repair systems?

RSVP and Physics

Unistochastic emergence. Under what conditions does RSVP transport generate a unistochastic transition kernel? The current conjecture (fibers admitting a $\mathrm{U}(n)$ action preserved by the dynamics) needs to be verified for the specific transport operators that arise in RSVP cosmology.

Ghost-parity in the spin-2 sector. The most pressing open problem in the physical applications is whether the ghost-parity symmetry of the Turok–Bateman Krein-space formulation of quadratic gravity extends to the spin-2 sector. If it does, the Derived Interface Conjecture gains strong support. If it does not, the conjecture must be revised.

RSVP field equations and general relativity. Under what conditions do the RSVP field equations reduce to the Einstein field equations in an appropriate limit? The leading-order admissibility-flow expansion is conjectured to yield general relativity, but the precise conditions need to be established.

Coordination and Collective Dynamics

HYDRA colimit existence. Under what conditions does the categorical colimit $\mathcal{A}_{\text{coll}} = \text{colim}_i \mathcal{A}_i$ exist and is non-empty? This is the formal question of when coordination is possible at all.

Phase transitions in collective admissibility. The transition from incoherence to synchronization in the Kuramoto model is a phase transition in the collective admissibility geometry. What are the analogous phase transitions in institutional systems, and can they be predicted from the geometry of the individual admissibility manifolds and the interaction maps?

30.2. Physical Open Problems

RSVP and the cosmological constant. The gravitational entropy program (Boyle and Turok) predicts that the observed value of the cosmological constant is typical for a universe with the observed entropy budget. Can this prediction be made quantitative from within the RSVP framework?

Dark matter and the CPT universe. The CPT-symmetric universe predicts that the stable right-handed neutrino has mass $M \approx 5 \times 10^8$ GeV to account for the observed dark matter density. This prediction will be testable by future large-scale structure surveys.

Primordial gravitational waves. The CPT-symmetric universe predicts no primordial tensor modes ($r = 0$ at leading order). Future experiments (CMB-S4, LiteBIRD) with sensitivity $r \lesssim 0.003$ will test this prediction.

Asymptotic freedom and the Higgs as a composite. The dimension-zero scalar program (Boyle and Turok) predicts that the Higgs is composite, with mass set by the dynamical scale of the asymptotically free conformal sector. The mechanism is analogous to dimensional transmutation in QCD. What is the predicted Higgs compositeness scale?

30.3. Conceptual Open Problems

The boundary between repair and learning. When does a system repair a collapsed distinction (restoration of what was) versus learn a new distinction (creation of something new)? The formal distinction between repair and learning in the distinction geometry framework has not been made precise.

Distinction creation and the origin of novelty. The framework accounts for the preservation, collapse, repair, and reorganization of distinctions. It does not yet give a full account of the creation of genuinely new distinctions: distinctions for which no prior distinction served as a template. This is the formal version of the problem of genuine novelty.

The geometry of institutional failure. The framework identifies institutional opacity, goal drift, and Goodhart effects as instances of hidden admissibility directions and distinction collapse. But it does not yet give a quantitative theory of institutional failure: at what rate do institutions drift, and what determines the rate?

Consciousness as distinction maintenance. If consciousness is the capacity to maintain distinctions — to be aware of what is different from what — then the framework provides the beginning of a formal account of consciousness. But the relationship between distinction maintenance and phenomenal experience is not yet worked out. This is the hardest problem in the program.

Appendices

Appendix A

Notation Reference

Symbol	Meaning
$(X, \sim)$	Distinction space: set X with equivalence relation $\sim$
$[x]_\sim$	Equivalence class of x
$X/\sim$	Quotient space of equivalence classes
$\Pi : M \to O$	Projection from history space to observable space
$F_o = \Pi^{-1}(o)$	Fiber over observable o
$\mathcal{A} \subseteq M$	Admissibility set
$d_{\mathcal{A}}(h)$	Admissibility distance of history h
$\Delta_R(h)$	Reconstruction defect of h under map R
$\delta(\Pi)$	Minimum worst-case reconstruction defect
$\mathcal{D}(P, Q)$	Distinction functional $=$ JSD$(P\|Q)$
$d_{\text{int}}(P, Q)$	Interface distance $=$ JSD$^{1/2}(P\|Q)$
$(\mathcal{M}, d_{\text{int}})$	Theory space
g^F_{ij}	Fisher information metric
$H(P)$	Shannon entropy of distribution P
M_{PQ}	Equal-weight mixture $\frac{1}{2}(P + Q)$
$Z(x)$	RSVP transport operator $\Phi(x)e^{\widehat{L}(x)}$
$\widehat{L}$	Transport generator $\mathbf{v} \cdot \nabla + \theta\widehat{T}$
T_{ij}	Observable transition matrix
$\mathcal{K} = \mathcal{K}_+ \oplus \mathcal{K}_-$	Krein space decomposition
$G = \Pi_+ - \Pi_-$	Ghost parity operator
$[u, v]_{\mathcal{K}}$	Krein inner product

Appendix B

Index of Frameworks

Admissibility Program

Distinguishability geometry and reachability-based ontology. Chapters 7–8, Parts III and VIII.

CLIO (Constraint-Layered Inference and Projection)

Formal theory of projection from internal to observable models. Chapter 5, Part II.

Coordination Geometry / HYDRA

Collective admissibility as categorical colimit; multi-agent coordination. Chapter 13, Part V.

Distinction Geometry

Jensen–Shannon interface metrics and the curvature of distinguishability. Chapter 27, Part IX.

Ecology of Distinctions

Network theory of distinction dependencies, conflicts, and cascades. Chapter 28, Part IX; also a 31-chapter standalone monograph.

MEM|8 (Wave-Stabilized Memory Architecture)

Formal framework for persistent memory in dynamic substrates. Chapter 11, Part IV. Authored by C. and A. Chenoweth.

Preference Fields

Geometry of value, preference, and decision over admissibility manifolds. Chapter 9, Part III.

Repair Theory

Formal theory of distinction restoration; repair as a cognitive primitive. Chapters 10–12, Part IV.

RSVP (Relativistic Scalar-Vector Plenum)

Field-theoretic cosmology grounded in constraint-first dynamics. Chapter 22, Part VIII.

Spherepop

Irreversible event calculus with Pop, Refuse, Bind, and Collapse operators. Chapter 18, Part VI. Implemented as a C interpreter.

TARTAN (Trajectory-Aware Recursive Tiling with Annotated Noise)

Computation on dynamic substrates with evolving admissibility manifolds. Chapter 20, Part VII.

Bibliography

[1] Flyxion, *The Ecology of Distinctions*, 31-chapter monograph on distinction theory, reachability, and admissibility geometry (in preparation).

[2] Flyxion, *Constraint, Projection, and Reachability* (CPR), 92-chapter LuaLaTeX monograph (in preparation).

[3] Flyxion, "Quadratic Gravity as a Derived Interface Theory: Admissibility Manifolds, Projection Geometry, and the Emergence of Observable Physics," preprint (2025).

[4] Flyxion, *Distinction Geometry: Projection, Reconstruction Defect, and the Interface-Native Metric* (Mathematical Core), preprint (2025).

[5] L. Boyle, K. Finn, and N. Turok, "CPT-symmetric universe," *Physical Review Letters* **121**, 251301 (2018).

[6] N. Turok and L. Boyle, "Gravitational entropy and the flatness, homogeneity and isotropy puzzles," *Physics Letters B* **849**, 138442 (2024).

[7] K. S. Stelle, "Renormalization of higher-derivative quantum gravity," *Physical Review D* **16**, 953–969 (1977).

[8] D. M. Endres and J. E. Schindelin, "A new metric for probability distributions," *IEEE Transactions on Information Theory* **49**(7), 1858–1860 (2003).

[9] S.-I. Amari and H. Nagaoka, *Methods of Information Geometry*, American Mathematical Society, Providence, 2000.

[10] T. Needham, *Visual Complex Analysis*, Oxford University Press, Oxford, 1997.

[11] J. A. Barandes, "The unistochastic quantum theory," *arXiv:2302.10778* [quant-ph] (2023).

[12] J. C. Scott, *Seeing Like a State*, Yale University Press, New Haven, 1998.

[13] T. S. Kuhn, *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, 1962.

[14] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM Journal of Research and Development* **5**(3), 183–191 (1961).

[15] E. T. Jaynes, *Probability Theory: The Logic of Science*, Cambridge University Press, Cambridge, 2003.

[16] K. M. Fant, *Logically Determined Design: Clockless System Design with NULL Convention Logic*, Wiley-Interscience, Hoboken, 2005.

[17] Y. Kuramoto, "Self-entrainment of a population of coupled non-linear oscillators," in *International Symposium on Mathematical Problems in Theoretical Physics*, Lecture Notes in Physics 39, Springer, Berlin, 1975, pp. 420–422.

[18] J. Kay and M. King, *Radical Uncertainty*, Bridge Street Press, London, 2020.