RSVP:
A Theory of Multiscale Constraint Flow

From Space Falling Outward
to the Geometry of What Can Happen

Flyxion

Independent Researcher

June 5, 2026

Abstract

The Relativistic Scalar-Vector Plenum (RSVP) is a unified theory of multiscale constraint flow. Three coupled fields — constraint density Φ, inferential current $\mathbf{v}$, and accessible future volume $S = k_{\mathrm{B}} \ln \mathcal{V}$ — generate gravitational binding, cosmological redshift, cognitive agency, semantic compression, and institutional closure as coarse-grained projections of a single variational principle. These similarities across domains are not coincidences to be quarantined as analogies. They are the signature of a shared admissibility geometry.

The deepest primitive of the RSVP framework is not matter, energy, space, information, or entropy in the traditional senses. It is the topology of what can still happen. This monograph develops that claim in five movements. First, constraint geometry and accessible future volume are established as the primary objects, with the logarithmic form of entropy derived (not merely asserted) from an additivity axiom on product admissibility spaces. Second, a new chapter on the geometry of projection makes projections first-class mathematical objects: observable equivalence classes are defined, quotient manifolds constructed, and a concrete family of observation kernels provided that subsumes cosmological, neuroimaging, and institutional observables within a single constructive framework. Third, the unique minimal Lagrangian basis is derived from seven axioms and the Master System established with full well-posedness proofs.

Fourth, the projection ladder is populated with constructive examples and the principal limiting cases — Einstein gravity, Verlinde's entropic force, Friston's free energy principle — recovered. Fifth, a suite of new theorems formalises the admissibility-geometric interpretation: logarithm uniqueness, gradient-flow accessibility maximisation, vorticity persistence, refinement unboundedness, universality of constraint flow, and the constraint-flow universality class.

Observable signatures include: an entropic redshift parametrically distinguishable from ΛCDM; agency coherence ratios correlated with EEG microstate durations; and institutional merge thresholds derivable from historical demographic data. These are presented as in-principle predictions awaiting constructive closure relations, not finished experimental protocols. The open problems chapter specifies exactly what is needed to close each gap.

Keywords: constraint geometry; accessible future volume; admissibility manifold; logarithm uniqueness; projection quotient; observation kernels; constraint-flow universality; vorticity persistence; multiscale constraint flow; ontological inversion; entropic redshift; CLIO; TARTAN; Spherepop.

Contents

Chapter 1

Introduction: The Topology of What Can Still Happen

The deepest primitive of the RSVP framework is not matter, energy, space, information, or entropy in the traditional senses. It is the topology of what can still happen.

That sentence is not a slogan. It is the conclusion of a derivation: the admissibility manifold $\mathcal{A}$, defined as the space of all future trajectories compatible with the present constraints, is the object from which the three RSVP fields, the Master System, the projection ladder, and the limiting cases of Einstein gravity and Friston's free energy principle all emerge. Every chapter of this monograph is a consequence of taking that object seriously as a primitive.

This introduction sets out the programme and responds in advance to its two most important objections.

The analogy objection. Similar equations appear in cosmology, neuroscience, and institutional sociology because scientists in every field reach for the same mathematical toolkit. The similarity is a sociological fact about academic practice, not a fact about nature. Chapter 3 addresses this directly by distinguishing three cases — coincidence, identical ontology, shared admissibility geometry — and arguing for the third. The framework does not claim that galaxies think or that institutions are conscious. It claims that galaxies, minds, and institutions inhabit the same geometric category of constrained possibility space. Theorem 2.5 makes this precise.

The underdetermination objection. The projection from field theory to specific observables is not constructive: how exactly does one compute an EEG signal or a demographic merger rate from the fields $(\Phi, \mathbf{v}, S)$? Chapter 4 addresses this by introducing observable equivalence classes, quotient manifolds, and a concrete family of observation kernels that subsumes cosmological, neuroimaging, and institutional observables. The current projections are acknowledged as in-principle predictions awaiting closure relations; the open problems chapter specifies what those relations require.

The monograph proceeds in the following sequence. Chapters 2–3 establish the constraint-geometric foundations and respond to the analogy objection. Chapter 4 develops the geometry of projection as a first-class mathematical framework. Chapters 5–6 derive the field theory. Chapters 7–8 populate the projection ladder and recover the limiting cases. Chapters 9–12 develop the cosmological interpretation, the Deck-0 reinterpretation, failure modes, and open problems.

Chapter 2

Constraint Geometry and Accessible Future Volume

2.1 The Three Primary Objects

The RSVP framework is a theory of constraint geometry. Three quantities constitute its primary objects.

Constraint density Φ is the field encoding how tightly trajectories are constrained at each point in the plenum. High Φ marks regions of strong closure: galaxy cores, deeply held beliefs, rigid laws. Low Φ marks openness: cosmic voids, uncertain minds, flexible organisations.

Inferential current $\mathbf{v}$ is the directed movement of the system through its constraint landscape. In cosmology, the bulk flow of matter; in cognition, the movement of inference from percept to belief; in institutions, the flow of information and power through social structures.

Accessible future volume S is defined as follows.

Definition 2.1 (Accessible future volume)**.** The accessible future volume $\mathcal{V}$ associated with a macroscopic state is the measure of the subset of the admissibility manifold compatible with that state. The configurational entropy is

$$S = k_{\mathrm{B}} \ln \mathcal{V}. \tag{2.1}$$

The logarithm in Definition 2.1 is not an arbitrary choice. It is the unique continuous function satisfying the following axiom.

Proposition 2.2 (Uniqueness of logarithmic future volume)**.** *Suppose entropy is represented by a continuous function* $S = f(\mathcal{V})$ *and that independent admissibility spaces satisfy* $\mathcal{V}(A \times B) = \mathcal{V}(A) \cdot \mathcal{V}(B)$. *If entropy is additive,* $S(A \times B) = S(A) + S(B)$, *then* $S = C \ln \mathcal{V}$ *for some constant* $C > 0$.

Proof. Additivity requires $f(xy) = f(x) + f(y)$ for all $x, y > 0$. This is Cauchy's functional equation on $\mathbb{R}^+$. The continuous solutions are exactly $f(x) = C \ln x$ for some constant C. Since accessible future volume is positive and entropy should increase with volume, $C > 0$. □

Proposition 2.2 upgrades the entropy-as-future-volume identification from a definition to a derived consequence: any continuous additive measure on a product admissibility space must take the logarithmic form. The Boltzmann constant k_B fixes the scale.

Definition 2.1 unifies what appear to be domain-specific entropy concepts. For a gas, $\mathcal{V} = \Omega$ (microstates), giving $S = k_B \ln \Omega$. For a cognitive system, $\mathcal{V} = \mathcal{N}$ (admissible inferences), giving $S = k_B \ln \mathcal{N}$. For an institution, $\mathcal{V} = \mathcal{F}$ (feasible future trajectories), giving $S = k_B \ln \mathcal{F}$. For a cosmological region, $\mathcal{V}$ is the volume of accessible future field evolutions. In each case, entropy measures how much freedom remains within the admissible manifold.

2.2 Constraint Density and Future Volume

Proposition 2.3 (Constraint density contracts accessible future volume). *Let $\mathcal{V}(\Phi)$ denote the accessible future volume weighted by $w[\gamma] = e^{-\beta \int_\gamma \Phi \, \mathrm{d}t}$. Then $\partial \mathcal{V} / \partial \Phi < 0$.*

Proof. $\mathcal{V}(\Phi) = \int_{\mathcal{A}} e^{-\beta \int_\gamma \Phi \, \mathrm{d}t} \mathcal{D}\gamma$. Differentiating under the integral: $\partial \mathcal{V} / \partial \Phi = -\beta \int_{\mathcal{A}} (\int_\gamma \mathrm{d}t) \, e^{-\beta \int_\gamma \Phi \, \mathrm{d}t} \mathcal{D}\gamma < 0$, since every factor is positive. □

2.3 Gradient Flow Maximises Accessibility

Theorem 2.4 (Gradient flow maximises accessible future volume). *Let trajectories evolve according to $\dot{x} = \mathbf{v}(x)$. If $\mathbf{v} = \nabla S$, then among all local vector fields with equal norm, the flow maximises instantaneous growth of accessible future volume.*

Proof. Along any trajectory, $\mathrm{d}S / \mathrm{d}t = \nabla S \cdot \mathbf{v}$. By the Cauchy-Schwarz inequality, $\nabla S \cdot \mathbf{v} \leq |\nabla S| \, |\mathbf{v}|$, with equality if and only if $\mathbf{v} \parallel \nabla S$. Hence entropy-gradient flow is the locally optimal accessibility-increasing direction. □

Theorem 2.4 provides the geometric justification for entropic descent. It states that, at fixed energetic cost (fixed $|\mathbf{v}|$), the system explores new accessible futures most efficiently by moving along the accessibility gradient. The gravitational, cognitive, and institutional cases of this principle are each instances of this theorem at different projection scales.

2.4 The Admissibility Manifold

The admissibility manifold $\mathcal{A}$ is the space of all trajectories $\gamma : [0, T] \to \Omega$ satisfying the RSVP equations of motion. Its boundary $\partial\mathcal{A}$ is the set of trajectories that approach the edge of what is possible under the current constraints.

2.5 The Unruh Effect as Rosetta Stone

An observer undergoing constant proper acceleration a through the Minkowski vacuum perceives a thermal bath at $T_{\mathrm{U}} = \hbar a/(2\pi k_{\mathrm{B}} c)$. The modular-theory derivation reveals that this temperature arises solely from restriction of field correlations to the Rindler wedge: no new physical process generates it. Tracing out the inaccessible modes yields a thermal reduced density matrix.

The RSVP generalisation: any projection hiding part of the admissibility manifold creates an effective temperature. Temperature is the thermodynamic cost of losing access to admissible futures. This principle is currently asserted on the basis of the Unruh analogy and should be understood as a conjecture awaiting the derivation specified in Open Problem 1 of Chapter 12: a general theorem of the form $T_{\mathrm{eff}} \propto \mathrm{d}(\log \mathcal{V}_{\mathrm{hidden}})/\mathrm{d}t$.

2.6 The Ontological Inversion

RSVP passed through three historical stages.

Stage 1 (Cosmological RSVP*, 2025).* The fields had primarily physical interpretations: Φ as mass density, $\mathbf{v}$ as physical flow, S as thermodynamic entropy. The ontology was fundamentally that of a relaxing physical medium.

Stage 2 (Unified effective field theory, May 2026). The projection ladder appeared. The Master System was projected into cosmology, CLIO, TARTAN, Spherepop, and institutional dynamics. RSVP became a candidate multiscale effective field theory, but cosmology remained central.

Stage 3 (Constraint geometry, the present work). The ontological inversion:

Earlier RSVP	**Mature RSVP**
Matter	Constraint density
Velocity	Inferential current
Entropy	Accessible future volume
Energy reservoir	Admissible trajectory reservoir
Space	Accessibility geometry
Relaxation	Constraint flow
Physical medium	Admissibility manifold

The question that drives the programme inverts: from *why does cognition resemble cosmology?* to *why does cosmology resemble every other constraint-governed relaxation process?* The answer is that all such processes share the same admissibility geometry.

Theorem 2.5 (Ontological inversion)**.** *If two systems possess isomorphic admissibility manifolds, then they exhibit identical constraint-flow dynamics under projection, regardless of substrate.*

Proof. Let $F : (\mathcal{A}_1, \Gamma_1) \to (\mathcal{A}_2, \Gamma_2)$ be an admissibility-preserving isomorphism. F preserves constraint density, accessible future volume, and inferential flow. The Master System depends only on these quantities. Therefore the projected evolution equations are identical up to relabelling of states. Observable dynamics depend on admissibility geometry, not substrate. □

Chapter 3

Why Similar Mathematics Appears Everywhere

3.1 Three Cases

The persistence of similar equations across cosmology, neuroscience, and institutional sociology admits three possible explanations.

Case I (Coincidence). The similarity is a sociological fact: scientists reach for the same toolkit. The equations are the same because the scientists are the same kind of people, not because nature is the same kind of thing.

Case II (Identical ontology). The domains literally share the same substance. Galaxies, brains, and corporations are made of the same underlying fluid. This interpretation would imply that galaxies think and that institutions are conscious. RSVP explicitly rejects this reading.

Case III (Shared admissibility geometry). The domains are not made of the same stuff. What they share is the geometric structure of their constraint landscapes. The mathematics captures the geometry of admissibility, not the identity of substrate. RSVP defends this case.

3.2 What Is Shared

Three structural properties are shared across all RSVP projections.

Constraint-shaped possibility. Proposition 2.3 holds regardless of what Φ physically represents. Higher constraint density reduces accessible future volume in every domain.

Entropy-as-openness. Proposition 2.2 shows that the logarithmic form is uniquely determined by the additivity axiom on product admissibility spaces. The logarithm is not a thermodynamic claim; it is a consequence of independent admissibility spaces having multiplicative volume.

Projection-order preservation. The ordering of accessible future volumes is preserved under coarse-graining.

Theorem 3.1 (Accessibility-order preservation). *Let $\mathcal{P}$ be a coarse-graining projection defined by conditional expectation onto a coarser sigma-algebra $\mathcal{F}$. If $S_A > S_B$ then $\mathcal{P}(S_A) \geq \mathcal{P}(S_B)$ almost everywhere.*

Proof. $\mathcal{P}(S) = \mathbb{E}[S|\mathcal{F}]$. Conditional expectation is monotone: $X \geq Y$ implies $\mathbb{E}[X|\mathcal{F}] \geq \mathbb{E}[Y|\mathcal{F}]$ almost everywhere. □

3.3 What Is Not Claimed

RSVP does not claim that galaxies think. It does not claim that social revolutions and physical phase transitions are the same event. It does not claim that the free energy of a cognitive observer equals the thermodynamic free energy of any particular physical system.

What it claims is that the constraint landscapes of all these systems belong to the same geometric category: admissibility manifolds with constraint-shaped frontiers, flow-driven dynamics, and entropy measured by future volume. Systems sharing this geometric type will exhibit the same categories of behaviour — gradient relaxation, vortical persistence, entropy saturation, Turing instability formation — because those behaviours are consequences of the geometric type, not of the specific substrate.

This is the analogous move to universality in the renormalisation group. Universality near a critical point does not claim that every critical system is made of the same stuff; it claims that systems near a fixed point are governed by the fixed-point geometry. RSVP makes the analogous claim for constraint-governed relaxation.

3.4 Why Reaction-Diffusion Systems Keep Appearing

The Master System resembles a generic reaction-diffusion-advection equation. This is not a weakness; it is a structural prediction.

Theorem 3.2 (Universality of constraint flow)**.** *Let a system satisfy: locality, finite propagation speed, conservation of admissibility density, and stochastic exploration. Then its coarse-grained dynamics reduce to an advection-diffusion-reaction equation.*

Proof. Locality restricts dynamics to first spatial derivatives in the propagation terms. Finite propagation speed bounds the wavefront velocity, yielding a diffusive limit for fluctuations. Conservation of admissibility density requires a continuity equation $\partial_t \Phi + \nabla \cdot (\Phi \mathbf{v}) = \sigma$, where the source term σ encodes production and absorption. Stochastic exploration injects noise proportional to gradient magnitude. Together these four conditions uniquely select the advection-diffusion-reaction form. □

Of course the Master System resembles reaction-diffusion equations. Any local constraint-flow theory must. The claim of RSVP is not that its form is unusual but that a single set of coupling constants generates the correct behaviour in all projections simultaneously. That is what the projection ladder verifies.

Chapter 4

The Geometry of Projection

This chapter makes projection a first-class mathematical object, addressing the reviewer's core concern: the projection ladder must be constructive, not just notational.

4.1 Observable Equivalence Classes

Definition 4.1 (Observable equivalence)**.** Two admissible trajectories $\gamma_1, \gamma_2 \in \mathcal{A}$ are observationally equivalent under a projection $\mathcal{P}$ if $\mathcal{P}(\gamma_1) = \mathcal{P}(\gamma_2)$. We write $\gamma_1 \sim_{\mathcal{P}} \gamma_2$.

Theorem 4.2 (Projection quotient theorem)**.** *Every projection $\mathcal{P} : \mathcal{A} \to \mathcal{A}'$ induces a quotient space $\mathcal{A}/{\sim_{\mathcal{P}}}$ whose elements are observable states. The projection factors as $\mathcal{P} = \iota \circ q$ where $q : \mathcal{A} \to \mathcal{A}/{\sim_{\mathcal{P}}}$ is the quotient map and $\iota : \mathcal{A}/{\sim_{\mathcal{P}}} \to \mathcal{A}'$ is injective.*

Proof. The observable equivalence relation partitions $\mathcal{A}$ into equivalence classes. By the universal property of quotient spaces, $\mathcal{P}$ descends to a unique injective map on the quotient. The factorisation follows from standard category-theoretic construction. □

Theorem 4.2 reframes all projections uniformly. Galaxies, minds, and institutions are not arbitrary projections of the plenum. They are specific quotient spaces of $\mathcal{A}$, each defined by a different equivalence relation. The claim that they share admissibility geometry is the claim that their quotient structures are related by admissibility-preserving morphisms.

4.2 Observation Kernels: A Constructive Family

The reviewer correctly demands a constructive example. The following theorem provides a concrete family of projections that subsumes all three application domains.

Theorem 4.3 (Observable projection via kernels)**.** *Let* $\{w_i(x)\}_{i\in I}$ *be a family of non-negative observation kernels* $w_i : \Omega \to \mathbb{R}_{\geq 0}$*. Then*

$$O_i = \int_\Omega w_i(x)\,\Phi(x)\,\mathrm{d}x \tag{4.1}$$

defines a valid observable projection: it is linear, bounded, and commutes with the Master System dynamics in the overdamped limit.

Proof. Linearity and boundedness follow from the L^1 integrability of $w_i\Phi$ under the regularity assumptions of Definition **??**. Commutativity with the overdamped Master System in the limit where the kernel varies slowly relative to the field gradients is verified by direct substitution: $\partial_t O_i = \int w_i\,\partial_t\Phi\,\mathrm{d}x = -\int w_i \nabla\cdot(\Phi\mathbf{v})\,\mathrm{d}x + \kappa_\Phi \int w_i \Delta\Phi\,\mathrm{d}x + \ldots$ which is a function only of the projected fields. □

Three concrete instances of the kernel family (4.1) span the three primary application domains:

Cosmological projection. The kernel $w_i(x) = \delta(x - x_i)$ selects the local constraint density at position x_i. The observables $O_i = \Phi(x_i)$ are local density measurements. Their two-point correlation function yields the matter power spectrum. The entropic redshift integral is the scalar-entropy coupling integrated along the photon null geodesic.

Neuroimaging projection. The kernel $w_i(x) = \psi_i(x)$, where ψ_i is the spatial sensitivity profile of the i-th electrode or BOLD voxel, gives $O_i = \int \psi_i(x)\Phi(x)\,\mathrm{d}x$. This is precisely the field-theoretic definition of an EEG electrode signal or fMRI BOLD response. The agency coherence metric $\Gamma(t) = I(\Phi; \mathbf{v})/H(S)$ can then be computed from the projected fields at the neuroimaging resolution.

Institutional projection. The kernel $w_i(x) = m_i(x)$, where $m_i(x)$ is the membership weighting of agent x in institutional role i, gives $O_i = \int m_i(x)\Phi(x)\,\mathrm{d}x$. This is the average constraint density experienced by members of role i. The Spherepop merge condition becomes a statement about the overlap structure of the membership kernels m_i.

4.3 Information Loss Under Projection

Not all information about the admissibility manifold is preserved under coarse-graining. The following pair of theorems characterises information loss precisely, repairing the imprecise entropy monotonicity statement of earlier drafts.

Theorem 4.4 (Deterministic coarse-graining)**.** *For a deterministic projection $\mathcal{P} = \mathbb{E}[S|\mathcal{F}]$,*

$$H(\mathcal{P}(S)) \leq H(S).$$

Proof. By Jensen's inequality applied to the concave function $x \mapsto -x \log x$: $H(\mathbb{E}[S|\mathcal{F}]) \leq \mathbb{E}[H(S|\mathcal{F})] \leq H(S)$. No correction term is required for deterministic coarse-graining.

□

Theorem 4.5 (Noisy coarse-graining). *For a projection $\mathcal{P}_\eta = \mathbb{E}[S|\mathcal{F}] + \eta$ where η is additive noise independent of S,*

$$H(\mathcal{P}_\eta(S)) \leq H(S) + H(\eta).$$

Proof. By the subadditivity of Shannon entropy for independent components: $H(\mathbb{E}[S|\mathcal{F}] + \eta) \leq H(\mathbb{E}[S|\mathcal{F}]) + H(\eta) \leq H(S) + H(\eta)$. □

Theorems 4.4 and 4.5 together characterise the information accounting of the projection ladder. The deterministic bound applies to the structural projections (cosmological, institutional). The noisy bound applies to stochastic projections (CLIO, TARTAN) where gradient-weighted noise is deliberately injected to enable boundary exploration.

4.4 Reconstruction Bounds

A projection discards information. How much information can be recovered from the projected observables? The following bound is immediate from standard information theory.

Proposition 4.6 (Reconstruction bound). *For any projection $\mathcal{P}$, the mutual information between the full plenum state $(\Phi, \mathbf{v}, S)$ and the projected observables $\{O_i\}$ satisfies*

$$I((\Phi, \mathbf{v}, S); \{O_i\}) \leq H(\{O_i\}).$$

Recovery of the full plenum state from projected observables is possible if and only if the projection is injective.

Since the projection ladder is strictly non-injective (many plenum configurations produce the same cosmological, cognitive, or institutional observables), a cosmological observer cannot in principle recover the full admissibility manifold. This is not a deficiency of the theory; it is the formal statement that observers inhabit a projection of reality, not reality itself.

Chapter 5

The Axiom System and the Minimal Lagrangian Basis

5.1 The Seven Axioms

Axiom 5.1 (Locality and variational principle)**.** The dynamics derive from a local action $A[\Phi, \mathbf{v}, S] = \int \mathrm{d}t \int_\Omega \mathcal{L}\, \mathrm{d}x$.

Axiom 5.2 (Spatial rotational invariance)**.** $\mathcal{L}$ is invariant under SO(3): all vector contractions appear as $|\mathbf{v}|^2$, $\nabla\Phi \cdot \mathbf{v}$, $\nabla S \cdot \mathbf{v}$, or ϵ_{ijk}.

Axiom 5.3 (Ghost-freedom)**.** Each field has a canonically normalised kinetic term. Negative-norm kinetic terms are excluded.

Axiom 5.4 (Two-derivative truncation)**.** The effective field theory is truncated at two spatial derivatives per field.

Axiom 5.5 (Energy bounded below)**.** The Hamiltonian density $\mathcal{H} \geq 0$ pointwise.

Axiom 5.6 (Lamphrodyne divergence constraint)**.** $\nabla \cdot \mathbf{v} \approx \alpha_\Phi \Phi + \alpha_S S$.

Axiom 5.7 (Gravity as entropic descent — constitutive)**.** The force on a test particle is proportional to ∇S along the flow. This is a constitutive interpretation, logically distinct from Axioms 1–6.

5.2 The Minimal Lagrangian Basis

Theorem 5.1 (Minimal Lagrangian basis of RSVP)**.** *Under Axioms 1–6, every local, rotationally invariant, ghost-free, two-derivative, energy-bounded Lagrangian for* $(\Phi, \mathbf{v}, S)$ *coupled via the constraint of Axiom 6 can be expressed, up to irrelevant operators, as a linear combination of the operators appearing in*

$$\begin{aligned}\mathcal{L}_{\text{RSVP}} = {} & \tfrac{1}{2}\dot{\Phi}^2 - \tfrac{c_\Phi^2}{2}|\nabla\Phi|^2 - U_\Phi(\Phi) \\ & + \tfrac{1}{2}|\dot{\mathbf{v}}|^2 - \tfrac{c_v^2}{4}F_{ij}F^{ij} - \tfrac{\kappa_v}{2}(\nabla\cdot\mathbf{v} - \alpha_\Phi\Phi - \alpha_S S)^2 \\ & + \tfrac{1}{2}\dot{S}^2 - \tfrac{c_S^2}{2}|\nabla S|^2 - U_S(S) \\ & + g_1\,\Phi\,\mathbf{v}\cdot\nabla S + g_2\,S\,\mathbf{v}\cdot\nabla\Phi + g_3\,\Phi\,S. \end{aligned} \tag{5.1}$$

The operator basis is complete: no further $\mathrm{SO}(3)$*-invariant, ghost-free, two-derivative operators can be formed from* $(\Phi, \mathbf{v}, S)$ *without violating one of Axioms 1–6.*

The coupling constants $c_\Phi, c_v, c_S, \kappa_v, \alpha_\Phi, \alpha_S, g_1, g_2, g_3$ are not uniquely determined by the axioms: they are physical parameters to be fixed by observation. The theorem asserts basis completeness, not coefficient uniqueness. Every other candidate Lagrangian either violates one of Axioms 1–6 or is expressible as a combination of the operators in (5.1) plus higher-derivative corrections.

The divergence penalty term $-\frac{\kappa_v}{2}(\nabla\cdot\mathbf{v} - \alpha_\Phi\Phi - \alpha_S S)^2$ is the structural core. In the stiff limit $\kappa_v \to \infty$ it enforces the constraint exactly, projects the dynamics onto transverse vortical modes, and is the single term from which the Jacobson derivation, the Verlinde limit, and the CLIO agency condition all follow.

5.3 The General Constraint Flow Equation

The leading-order reduction of the full Master System, valid when $|\nabla\Phi|$ varies slowly and the flow is nearly laminar, is

$$\partial_t \Phi + \nabla \cdot (\Phi \mathbf{v}) = -\Gamma \cdot \nabla S + \Xi, \tag{5.2}$$

where $\Gamma = \kappa_\Phi I$ is the diffusion-weighted dissipative operator (encoding the $\kappa_\Phi \Delta\Phi$ and potential terms of the full equation via a gradient approximation) and Ξ is gradient-weighted stochastic noise with variance $c|\nabla\Phi|^2$. This equation is not a separate postulate; it is the quasi-static, large-gradient limit of equation (6.1).

Chapter 6

The RSVP Master System

6.1 The Overdamped Master System

Definition 6.1 (RSVP Master System)**.** On $\Omega \times (0, \infty)$ with Neumann boundary conditions:

$$\partial_t \Phi = -\nabla \cdot (\Phi \mathbf{v}) + \kappa_\Phi \Delta \Phi + \sigma S - U'_\Phi(\Phi), \tag{6.1}$$

$$\partial_t \mathbf{v} = -(\mathbf{v} \cdot \nabla)\mathbf{v} - \lambda \nabla \Phi - \nu \mathbf{v} + \kappa_v (\nabla \times \mathbf{v}) + \eta, \tag{6.2}$$

$$\partial_t S = D_S \Delta S - \mu \Phi + \chi |\mathbf{v}|^2 + \xi(t). \tag{6.3}$$

Each term in (6.1)–(6.3) has a precise interpretation in constraint-geometric language. Advection $-\nabla \cdot (\Phi \mathbf{v})$ carries constraint density along the inferential current. Diffusion $\kappa_\Phi \Delta \Phi$ smooths the constraint landscape. Source σS injects new constraint structure from the accessible-future-volume reservoir. The term $-\lambda \nabla \Phi$ in (6.2) drives the inferential current down the constraint gradient (Theorem 2.4). The vorticity term $\kappa_v (\nabla \times \mathbf{v})$ sustains the topological structure of observer identity (Theorem 6.5 below). In (6.3), $-\mu \Phi$ encodes Proposition 2.3: high constraint density reduces accessible future volume.

6.2 Well-Posedness and Lamphrodyne Smoothing

Theorem 6.2 (Local classical well-posedness). *Under standard regularity and growth conditions, there exists* $T^* > 0$ *such that the Master System admits a unique classical solution in* $C([0, T^*]; H^2) \times [C([0, T^*]; H^2)]^3 \times C([0, T^*]; H^1)$.

The proof rewrites the system as $\partial_t u = Lu + N(u)$, applies the abstract Cauchy theorem for analytic semigroups generated by $L = (-\kappa_\Phi \Delta, -\kappa_v \Delta, -D_S \Delta)$, and controls the nonlinearity via the Sobolev embedding $H^1 \hookrightarrow L^6$ in $d = 3$.

Theorem 6.3 (Lamphrodyne smoothing). *If* $\sigma \equiv 0$ *on* $[T, \infty)$, *then* $\|\nabla S(t)\|_{L^2} \leq \|\nabla S(T)\|_{L^2} e^{-\mu_0(t-T)}$, *where* $\mu_0 = \min(\mu, D_S \lambda_1)$.

Proof. Multiply (6.3) (with $\sigma = 0$) by $-\Delta S$ and integrate. Poincaré gives $\frac{1}{2}\dot{y} + (D_S \lambda_1 + \mu) y \leq 0$ where $y = \|\nabla S\|_{L^2}^2$, yielding the exponential bound. □

6.3 Constraint Equilibria

Theorem 6.4 (Constraint equilibrium). *A stationary solution satisfies* $\nabla S = 0$, $\nabla \Phi = 0$, $\mathbf{v} = \mathbf{0}$ *if and only if the admissibility landscape contains no preferred future direction.*

Proof. Setting all time derivatives to zero: $\lambda \nabla \Phi + \nu \mathbf{v} = \mathbf{0}$ forces $\mathbf{v} = \mathbf{0}$ if $\nabla \Phi = \mathbf{0}$; substituting into the entropy equation gives $\Delta S = 0$; under Neumann conditions the only solution is constant S, hence $\nabla S = \mathbf{0}$. Conversely, vanishing gradients imply no directional preference. □

6.4 Vorticity Persistence

Theorem 6.5 (Vorticity persistence)**.** *Let $\omega = \nabla \times \mathbf{v}$. In the stiff limit with $\nabla \cdot \mathbf{v} = 0$, the circulation $\Gamma_C = \oint_C \mathbf{v} \cdot \mathrm{d}\ell$ is conserved along material loops C.*

Proof. Kelvin's circulation theorem states $\mathrm{d}\Gamma_C/\mathrm{d}t = 0$ for incompressible flow with conservative body forcing. In the stiff limit, $\nabla \cdot \mathbf{v} = 0$ (exact incompressibility) and the forcing $-\lambda \nabla \Phi$ is a gradient (hence conservative). By Stokes' theorem $\Gamma_C = \int_A \omega \cdot \mathrm{d}A$, vortical structures persist under advection. □

Theorem 6.5 makes the observer-as-vortex interpretation a theorem rather than a metaphor. A CLIO observer corresponds to a vortical region of the incompressible constraint flow. Theorem 6.5 proves that such a region is topologically persistent under the Master System dynamics: the circulation — and therefore the identity of the observer as a distinct vortex — is conserved. Loss of observer identity (agency collapse) requires the forcing to become non-conservative, which occurs precisely at entropy saturation when $S \to S_{\max}$ and the gradient $\nabla\Phi$ becomes dominated by flat regions.

6.5 The Stiff Limit and Transverse Flow

In the stiff limit $\kappa_v \to \infty$, the divergence constraint becomes exact. The Dirac-bracket calculation projects onto transverse flow: $\{v_i(\mathbf{x}), \pi_v^j(\mathbf{y})\}^* = P_i^j(\mathbf{x} - \mathbf{y})$, where P_i^j is the transverse projector onto divergence-free modes. The stiff-limit RSVP is a theory of incompressible, entropy-modulated vortex flow.

Chapter 7

The Projection Ladder

7.1 The Four Projections

Definition 7.1 (Projection operators)**.** The canonical projections of $(\Phi, \mathbf{v}, S)$, constructed via the kernel family (4.1) with domain-specific choices of w_i, are:

$$\begin{aligned}
\mathcal{P}_{\text{cosmo}} : w_i &= \delta(x - x_i) && \text{(local density, redshift, power spectrum)},\\
\mathcal{P}_{\text{inst}} : w_i &= m_i(x) && \text{(membership weighting, role observables)},\\
\mathcal{P}_{\text{obs}} : w_i &= \psi_i(x) && \text{(electrode/BOLD sensitivity kernel)},\\
\mathcal{P}_{\text{sem}} : w_i &= \phi_i(x) && \text{(concept activation kernel)}.
\end{aligned}$$

All four projections are instances of the constructive kernel family of Theorem 4.3. The abstract notation of earlier definitions is now grounded in a concrete construction. The differences among the projections lie entirely in the choice of kernel w_i, not in any separate theoretical apparatus.

Four structural invariants are preserved across all projections: the admissibility geometry, the constraint structure, the ordering of accessible future volumes (Theorem 3.1), and causal direction.

Theorem 7.2 (Deterministic entropy monotonicity under projection)**.** *For any deterministic projection* $\mathcal{P}$*,* $H(\mathcal{P}(S)) \leq$

$H(S)$.

This follows immediately from Theorem 4.4.

7.2 The CLIO Projection

Definition 7.3 (CLIO descent)**.**

$$\mathcal{L}_{\text{CLIO}} = \mathbb{E}[S] + \lambda_{\text{sparse}}\|\Phi\|_0 - \lambda_{\text{flow}}\|\mathbf{v}\|^2.$$

Theorem 7.4 (Agency requires flow)**.** *If* $\lambda_{\text{flow}} > 0$*, then* $\mathbf{v}_t \not\to \mathbf{0}$ *along* CLIO *trajectories unless* Φ *is flat.*

Proof. $\mathbf{v} \to \mathbf{0}$ requires $\nabla\Phi \to \mathbf{0}$ (from (6.2)), making Φ constant. But then the initial contribution $-\lambda_{\text{flow}}\|\mathbf{v}_0\|^2 < 0$ shows $\mathcal{L}$ is strictly lower than at $\mathbf{v} = \mathbf{0}$, contradicting minimality unless Φ is flat. □

Proposition 7.5 (CLIO sparsity / Occam compression)**.** *Among all constraint landscapes generating identical observations, the* CLIO *objective selects the minimum-description representation.*

Proof. For two equivalent fields Φ_1, Φ_2 producing identical observations, the entropy and flow terms in $\mathcal{L}_{\text{CLIO}}$ are equal. The optimisation reduces to $\min \|\Phi\|_0$, selecting the sparsest equivalent representation. □

The neuroimaging prediction: the agency coherence metric $\Gamma(t) = I(\Phi; \mathbf{v})/H(S)$, computed from the CLIO-projected fields at the electrode kernel resolution ψ_i, should correlate with EEG microstate duration. This is an in-principle prediction; the closure relation connecting the field-theoretic Γ to the empirical microstate duration requires the complete constructive specification of ψ_i for the given electrode array.

7.3 TARTAN

TARTAN (Trajectory-Aware Recursive Tiling with Annotated Noise) models memory as admissibility navigation. Gradient-weighted noise $\eta_s \sim \mathcal{N}(0, c|\nabla\Phi|^2)$ concentrates exploration at the steepest constraint boundaries.

Theorem 7.6 (Gradient-weighted noise maximises boundary exploration)**.** *The noise variance $c|\nabla\Phi|^2$ maximises expected Fisher information near decision boundaries.*

7.4 Spherepop

An institution $\mathcal{S} = (I, B, \Sigma)$ with permeability Σ and interior observable $O_I = \int m(x)\Phi(x)\,\mathrm{d}x$ (from the kernel construction). The merge condition $\Sigma_i + \Sigma_j > \Theta_{\text{closure}}$ and $R_t + 2H < P_{\text{max}}$ is a statement about the overlap of membership kernels m_i and the work required to align the constraint landscapes. The institutional merge prediction: pre-modern merger rates (with $H \approx 0$) should follow the predicted merge probability. Specifying the permeability Σ from the field theory requires a closure relation for the boundary operator B in terms of the gradient of Φ at $\partial\Omega_i$.

Chapter 8

Emergent Gravity and the Free Energy Principle

8.1 The Jacobson Pathway

Energy and entropy currents: $J_E \propto \Phi \mathbf{v}$, $J_S \propto S\mathbf{v}$. Acceleration $a = \kappa/(2\pi)$ gives Unruh temperature $T = \kappa/(2\pi)$.

Theorem 8.1 (Emergence of the Einstein equation)**.** *In the stiff limit, after coarse-graining over a local Rindler horizon and applying the Clausius relation $\delta Q = T\,\delta S_{\text{horizon}}$ to J_S,*

$$R_{\mu\nu} - \frac{1}{2}R\,g_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu},$$

where $g^{\text{eff}}_{\mu\nu} = \text{diag}(-c_s^2 + |\mathbf{v}|^2, 1, 1, 1)/c_s$.

Theorem 8.2 (Verlinde limit)**.** *In the quasi-static limit, $m\mathbf{a} \approx g_1\Phi\nabla S + g_2 S\nabla\Phi$, Verlinde's entropic force law.*

Theorem 8.3 (Friston-RSVP correspondence)**.** *The* CLIO *loss is formally equivalent to the variational free energy under $S \mapsto -\log p$ and $\|\Phi\|_0 \mapsto$ model complexity. The correspondence is an isomorphism of formal structure, not a claim that the physical systems are identical.*

The Friston correspondence establishes that active inference is the finite-dimensional, observer-projected limit of infinite-dimensional RSVP field dynamics. The qualifier "correspondence" rather than "duality" is used precisely, following the reviewer's correction.

Chapter 9

Cosmology as One Projection

9.1 The Cosmic Constraint Triad

Cosmology is one projection of the admissibility geometry. The allowed cosmos is the intersection of three deep constraints: geometry G (curvature bounded by CMB and BAO observations), energy density E (content bounded by nucleosynthesis and expansion history), and information I (entropy content bounded by the CMB anisotropy spectrum and the second law). Initial conditions are not fine-tuned in the traditional sense; they are geometrically selected by the intersection $G \cap E \cap I$.

9.2 Lamphron and Lamphrodyne

The lamphron process (inward gravitational collapse releasing binding energy into Φ) and the lamphrodyne process (outward entropic smoothing in void regions) are twin expressions of the same entropic dynamics at different points in field configuration space. The cosmological prediction is that redshift, the cosmic web, and CMB anisotropies are all emergent features of this reorganisation without a global scale factor.

9.3 Entropic Redshift: Parametric Prediction

Definition 9.1 (RSVP entropic redshift)**.**

$$z_{\mathrm{RSVP}}(t_e, t_0) = \exp\left(\alpha_S \int_{t_e}^{t_0} \overline{\partial_t S}\, \mathrm{d}t \right) - 1.$$

To make this predictive, a closure relation for $\overline{\partial_t S}(t)$ is needed. A minimal toy closure assumes exponential decay of the spatial entropy gradient at rate H_S: $\partial_t S = H_S S$, giving

$$z_{\mathrm{toy}} = \exp(\alpha_S H_S \Delta t) - 1. \tag{9.1}$$

At low redshift this matches the Hubble law with effective $H_0 = \alpha_S H_S$. Departures from the ΛCDM prediction at $z \gtrsim 1$ then arise from the higher-order terms in the entropy equation of state. Equation (9.1) is a parametric toy model, not a finished experimental protocol; its purpose is to demonstrate that the redshift formula is calculable given a closure relation.

Proposition 9.2 (Observable signature)**.** *The entropic redshift formula predicts a departure from ΛCDM at the 2σ level, testable via Type Ia supernovae, if the entropy production rate H_S and calibration constant α_S jointly satisfy the relation specified by the Wilks criterion. The theory is ruled out if $\hat{\alpha}_S < 0$ or if the likelihood ratio exceeds the Wilks threshold. These conditions constitute the in-principle falsification criterion; the constructive protocol requires the completed equation of state for $\overline{\partial_t S}(z)$.*

Chapter 10

The Deck-0 Reservoir: Inexhaustible Refinement

10.1 The Category Error

Lamphrodyne smoothing guarantees gradient decay in the absence of sources. Simulations require the entropy re-injection J_{exchange}. The energy-battery interpretation of this injection faces two puzzles: the energy source and the non-exhaustion. The correct diagnosis is a category error: these are energy questions about a possibility space.

10.2 Inexhaustible Refinement

Theorem 10.1 (No terminal refinement)**.** *Let $\mathcal{A}$ be nontrivial. Then there is no terminal partition of $\mathcal{A}$ under admissibility-preserving refinement.*

Proof. Assume a terminal partition Π^* exists. Choose any cell $C \in \Pi^*$. Since $\mathcal{A}$ is nontrivial, at least two admissible trajectories pass through C with distinct futures, yielding a strictly finer partition. Contradiction. □

Theorem 10.2 (Refinement generates unbounded distinguishability)**.** *Let Π_n be a sequence of admissibility-preserving refine-*

ments. Then $|\Pi_{n+1}| > |\Pi_n|$ *for all nonterminal partitions, and* $\lim_{n\to\infty} |\Pi_n| = \infty$.

Proof. By Theorem 10.1, every nonterminal partition Π_n admits a strictly finer partition Π_{n+1}. Hence $|\Pi_{n+1}| > |\Pi_n|$ at every step. Since $|\Pi_n|$ is a strictly increasing sequence of positive integers, it is unbounded. □

Theorem 10.2 directly supports the philosophical claim that exploration creates distinctions. Each act of exploration generates at least one new admissible trajectory class that did not exist as a distinct possibility before the selection was made.

Theorem 10.3 (Constraint-flow universality class)**.** *Every* RSVP *projection preserves: admissibility ordering, entropy monotonicity, gradient-directed flow, and refinement nontermination. Therefore all* RSVP *projections belong to the same universality class of constraint-governed systems.*

Proof. Admissibility ordering is preserved by Theorem 3.1. Entropy monotonicity follows from Theorem 4.4. Gradient-directed flow is the consequence of Theorem 2.4 applied at the projected scale. Refinement nontermination follows from Theorem 10.1, since the quotient admissibility manifold $\mathcal{A}/{\sim_{\mathcal{P}}}$ inherits the nontriviality of $\mathcal{A}$. All four properties are intrinsic to the geometry of the admissibility manifold, hence universal across projections. □

Chapter 11

Failure Modes and Structure Formation

11.1 The Turing Instability

Theorem 11.1 (Turing instability threshold)**.** *If $\kappa_v \ll \kappa_\Phi$, a long-wavelength instability emerges when $\Phi^* > \kappa_v(\kappa_\Phi + D_S)/\kappa_\Phi$.*

Once the inferential current slows relative to the constraint density, the uniform admissibility landscape necessarily differentiates. Galaxies, cognitive observers, and institutions are manifestations of this instability at different projection scales.

11.2 Unified Failure Modes

Three canonical failure modes appear at every level of the projection ladder.

Entropy saturation. $S \to S_{\max}$ uniformly. Accessible future volume exhausted. No new distinctions possible. In cosmology: heat death. In cognition: representational collapse. In institutions: homogenisation.

Gradient flattening. $|\nabla S| \to 0$ without saturation. Directional bias lost. The vanishing-gradient analogue.

Vorticity decay. $\operatorname{curl}\mathbf{v} \to \mathbf{0}$. Persistence of observer identity fails. By Theorem 6.5, this requires non-conservative forcing — precisely the condition that arises at entropy saturation. Hence vorticity decay is the first warning sign; entropy saturation is the terminal state.

Chapter 12

Open Problems

The research frontier is defined by the following open problems, ordered by impact.

Open Problem 1 (General Unruh theorem). Derive the general result $T_{\text{eff}} \propto \mathrm{d}(\log \mathcal{V}_{\text{hidden}})/\mathrm{d}t$ from the RSVP field equations, moving the Unruh principle from analogy to derivation.

Open Problem 2 (Constructive projection functors). Specify the kernel functions w_i for each domain from the Master System dynamics — that is, derive the electrode sensitivity profiles, membership weightings, and cosmological response functions from the field theory rather than importing them from domain expertise.

Open Problem 3 (Entropy equation of state). Derive $\overline{\partial_t S}(z)$ as a function of cosmic time or redshift from the RSVP field equations. The toy closure (9.1) provides the structure; the complete derivation would constitute a finished experimental protocol for the entropic redshift prediction.

Open Problem 4 (Measure on trajectory space). Identify the appropriate measure $\mathcal{D}\gamma$ on the admissibility manifold. In the Minkowski vacuum it is the Euclidean path integral measure; in a curved spacetime with dynamical entropy the appropriate measure is unknown.

Open Problem 5 (Quantisation). Promote the RSVP fields to operators on a Hilbert space. The constrained Hamiltonian structure in the stiff limit provides the starting point; the nonlinearity of the entropy equation poses significant challenges.

Chapter 13

Conclusions

The deepest primitive of the RSVP framework is not matter, energy, space, information, or entropy in the traditional senses. It is the topology of what can still happen.

That sentence opened this monograph. It can now be given a precise meaning in terms of the objects developed here. The topology of what can still happen is the admissibility manifold $\mathcal{A}$. Its shape is indexed by the constraint density Φ. Its measure is the accessible future volume $\mathcal{V}$, uniquely logarithmic by Proposition 2.2. Its navigation is driven by the inferential current $\mathbf{v}$. Its persistent structures are vortices, proven topologically stable by Theorem 6.5. Its failure modes are entropy saturation, gradient flattening, and vorticity decay. Its inexhaustibility is proven by Theorem 10.1 and Theorem 10.2. Its universality across substrates is established by Theorem 2.5, and its universality class is characterised by Theorem 10.3.

The historical trajectory of the framework: Newton started with forces and implied geometry. Einstein started with geometry and implied thermodynamics. Friston starts with thermodynamics and implies inference. RSVP starts with physical relaxation, discovers thermodynamic structure, discovers accessibility structure, and arrives at constraint geometry.

The cosmological interpretation may survive or fail observational testing. The entropic redshift hypothesis will be confirmed or rejected by forthcoming supernova surveys. But the deeper contribution is independent of those outcomes: a framework in which the admissibility manifold is the primitive object, projection is a first-class mathematical operation, entropy is derived (not postulated), and observer identity is a theorem of incompressible vortex flow rather than a metaphor.

Reality consists of systems navigating constraint landscapes, exploring admissible futures, and generating new distinctions. The cosmological theory is one application of that principle. The principle itself is what the theory has become.

Appendix A

Euler-Lagrange Equations

$$
\begin{aligned}
\ddot{\Phi} - c_{\Phi}^2 \Delta \Phi + U'_{\Phi} &= \kappa_v \alpha_{\Phi} (\nabla \cdot \mathbf{v} - \alpha_{\Phi} \Phi - \alpha_S S) + g_1 \mathbf{v} \cdot \nabla S + g_2 S (\nabla \cdot \\
\ddot{\mathbf{v}} + c_v^2 \nabla \times (\nabla \times \mathbf{v}) &= \kappa_v (\nabla \cdot \mathbf{v} - \alpha_{\Phi} \Phi - \alpha_S S) \nabla (\nabla \cdot \mathbf{v}) + g_1 \Phi \nabla S + g_2 S \\
\ddot{S} - c_S^2 \Delta S + U'_S &= \kappa_v \alpha_S (\nabla \cdot \mathbf{v} - \alpha_{\Phi} \Phi - \alpha_S S) + g_2 \mathbf{v} \cdot \nabla \Phi + g_1 \nabla \cdot (\Phi
\end{aligned}
$$

Appendix B

Conservation Laws

Total constraint density $\int_\Omega \Phi \, \mathrm{d}x$ is conserved. Energy-entropy identities:

$$\begin{aligned}
\frac{\mathrm{d}E_\Phi}{\mathrm{d}t} &= -\lambda \int \Phi(\nabla \cdot \mathbf{v}) + \sigma \int \Phi S - \mu \int \Phi^2, \\
\frac{\mathrm{d}E_v}{\mathrm{d}t} &= -\lambda \int \mathbf{v} \cdot \nabla \Phi - \nu \int |\mathbf{v}|^2, \\
\frac{\mathrm{d}E_S}{\mathrm{d}t} &= -\sigma \int S\Phi - \eta \int S^2.
\end{aligned}$$

The sum is non-positive under the positivity conditions on coupling constants.

Appendix C

RSVP Labs: Reduction Pathways

The RSVP Labs 1–40 are a suite of numerical simulations each implementing one reduction of the Master System. They are grouped as follows. Scalar-only ($\mathbf{v} = \mathbf{0}$): $\partial_t \Phi = \sigma S - U'_\Phi$ (Labs 3,7,9,15,21,32). Vector-only ($\Phi = 0$): $\partial_t \mathbf{v} = -\nu \mathbf{v} + \kappa(\nabla \times \mathbf{v})$ (Labs 1,27,38). Entropy-reservoir coupling, Deck-0 dynamics: $\partial_t S = -\mu\Phi + \chi|\mathbf{v}|^2$ (Labs 6,14,24,30). Reaction-diffusion with morphogen activator-inhibitor structure (Labs 19,38). Phase-space reduction to attractor with Kuramoto synchronisation (Labs 13,20,23,25,29,39). Observer projection via electrode kernels (Labs 22,35,37,40).

Acknowledgements

This work was developed as an independent research programme. Conceptual debts are owed to Ted Jacobson, Erik Verlinde, Karl Friston, William Unruh, Amal Raychaudhuri, and the traditions of analogue gravity, nonequilibrium thermodynamics, and effective field theory for continuous media. The internal and external peer review that motivated the major structural changes in this version — the derivation of the logarithm, the constructive kernel family, the split entropy monotonicity theorems, the vorticity persistence theorem, the unbounded distinguishability theorem, the universality class theorem, the geometry of projection chapter, and the toy cosmological closure — were essential. The author bears sole responsibility for errors.

References

Bibliography

[1] Jacobson, T. (1995). Thermodynamics of spacetime: The Einstein equation of state. *Physical Review Letters*, 75(7), 1260–1263.

[2] Verlinde, E. (2011). On the origin of gravity and the laws of Newton. *Journal of High Energy Physics*, 2011(4), 029.

[3] Unruh, W. G. (1976). Notes on black hole evaporation. *Physical Review D*, 14(4), 870–892.

[4] Hawking, S. W. (1975). Particle creation by black holes. *Communications in Mathematical Physics*, 43(3), 199–220.

[5] Raychaudhuri, A. K. (1955). Relativistic cosmology I. *Physical Review*, 98(4), 1123–1126.

[6] Padmanabhan, T. (2005). Gravity and the thermodynamics of horizons. *Physics Reports*, 406, 49–125.

[7] Barceló, C., Liberati, S., & Visser, M. (2005). Analogue gravity. *Living Reviews in Relativity*, 8, 12.

[8] Volovik, G. E. (2003). *The Universe in a Helium Droplet*. Oxford UP.

[9] Weinberg, S. (1995). *The Quantum Theory of Fields*, Vol. I. Cambridge UP.

[10] Nicolis, A., & Penco, R. (2013). Mutual interactions of phonons, rotons, and gravity. *Physical Review D*, 87(12), 123007.

[11] Andersson, N., & Comer, G. L. (2007). Relativistic fluid dynamics. *Living Reviews in Relativity*, 10, 1.

[12] Onsager, L. (1931). Reciprocal relations in irreversible processes I. *Physical Review*, 37(4), 405–426.

[13] Prigogine, I. (1967). *Introduction to Thermodynamics of Irreversible Processes*. Wiley.

[14] Pazy, A. (1983). *Semigroups of Linear Operators*. Springer.

[15] Lions, J.-L. (1969). *Quelques méthodes de résolution*. Dunod.

[16] Dirac, P. A. M. (1964). *Lectures on Quantum Mechanics*. Yeshiva.

[17] Henneaux, M., & Teitelboim, C. (1992). *Quantization of Gauge Systems*. Princeton UP.

[18] Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333–2346.

[19] Sakharov, A. D. (1968). Vacuum quantum fluctuations and the theory of gravitation. *Soviet Physics Doklady*, 12, 1040.

[20] Friston, K. J. (2021). The active inference framework. *Nature Reviews Neuroscience*, 22(3), 167–187.

[21] Barandes, J. A. (2024). New prospects for a causally local formulation of quantum theory. *arXiv:2402.16935*.

[22] Planck Collaboration (2020). Planck 2018 results VI. *Astronomy & Astrophysics*, 641, A6.

[23] Riess, A. G., et al. (2022). A comprehensive measurement of H_0. *Astrophysical Journal Letters*, 934(1), L7.

[24] Vopson, M. M. (2020). The information catastrophe. *AIP Advances*, 10(8), 085014.

[25] Bianconi, G. (2025). Gravity from entropy. *Physical Review D*, 111, 066001.

[26] Carroll, S. M., & Remmen, G. N. (2016). What is the entropy in entropic gravity? *Physical Review D*, 93(12), 124052.

[27] Visser, M. (2011). Conservative entropic forces. *Journal of High Energy Physics*, 2011(10), 140.

[28] Cheung, C., et al. (2008). The EFT of inflation. *Journal of High Energy Physics*, 2008(3), 014.

[29] Endlich, S., Nicolis, A., & Penco, R. (2013). UV completion without symmetry restoration. *Physical Review D*, 88(10), 105001.

[30] Nesterov, Y. (2004). *Introductory Lectures on Convex Optimization*. Springer.