

Bayesian Statistics and Photonics

Failures of Traditional Models and Solutions
via the RSVP Framework
and the Spherepop/KES Calculus

Flyxion
Independent Researcher

Preface

This textbook undertakes a foundational reconstruction of Bayesian inference and quantum photonics. It begins from the claim that existing frameworks operate on the wrong primitive objects. Standard Bayesian inference takes probability distributions over predefined hypothesis spaces as its primary objects. Standard quantum optics takes Hilbert space vectors and the Born rule as its primary objects. Both are compressions of a richer underlying structure: the admissible trajectory field from which observable quantities and statistical distributions both emerge.

The book is organized around four structural ideas that recur at every level of the exposition. The first is *admissibility*: not all configurations are equally accessible at a given time, and this differential accessibility is the true primitive from which probability distributions emerge. The second is *projection*: the map from rich trajectory space to compressed representational manifold, whose fiber structure encodes the information lost by conventional models. The third is *entropy as accessibility*: the entropy field S measures future trajectory volume, not merely Shannon uncertainty. The fourth is *irreversible event structure*: the Spheredop/KES calculus treats events as primitive and derives states as equivalence classes of event histories, placing irreversibility at the foundation rather than deriving it from reversible dynamics.

Every chapter of this text is anchored to one or more of these four ideas. The reader who loses the thread can always return to them.

The book is organized as follows. Part 0 covers mathematical prerequisites: Hilbert spaces, manifolds, Fisher-Rao geometry, and rigorous Bayesian inference. Part I is diagnostic, identifying specific failures of orthodox models. Part II develops the RSVP framework as a field-theoretic replacement, beginning with admissibility geometry as the conceptual foundation. Part III develops the Spherepop/KES irreversible event calculus. Part IV unifies the two frameworks, compares them to existing programs, and catalogues experimental predictions. Part V develops extensions: projection failure, information geometry beyond Fisher-Rao, non-Markovian dynamics, photonics as constraint propagation, topological defects, thermodynamics, learning, and cosmological optics. Part VI closes with a unified accessibility calculus and open problems.

Epistemic status labels used throughout:

Theorem/Proposition Mathematically derived from stated assumptions.

Conjecture Believed true but not yet proved within this framework.

Ansatz A working assumption adopted for tractability.

Interpretation Physical or inferential reading of a formal result.

Prediction Empirically testable consequence, dependent on stated ansätze.

Flyxion, Independent Researcher

Notation and Symbols

Symbol	Name	Interpretation
Φ	Scalar field	Likelihood potential / salience landscape
\mathbf{v}	Vector field	Admissibility flow / information transport
S	Entropy field	Local accessibility density
(Φ, \mathbf{v}, S)	RSVP triple	The full RSVP field configuration
\mathcal{M}	Plenum manifold	Parameter / state space
g	Riemannian metric	Geometry of the plenum
X	Trajectory space	Space of all admissible histories
M	Representational manifold	Compressed model space
$\pi : X \rightarrow M$	Projection	Compression map
$\pi^{-1}(m)$	Fiber	Trajectories projecting to $m \in M$
$\mathcal{A}(x, t)$	Admissibility	Dynamic accessibility of x at time t
Ω_t	World-state	Constraint-weighted distribution
H_t	History state	Realized causal event sequence
Pop	Pop operator	Event to immediate consequence
Bind	Bind operator	Joins causally compatible events
Collapse	Collapse operator	Identifies compatible configurations
Refuse	Refuse operator	Filters events by consistency
$\mathcal{F}(e)$	KES functional	Selection weight for event e
π_{RSVP}	RSVP prior	Dynamically generated prior
ρ	Density operator	Quantum optical state
E_k	POVM element	Measurement operator
p_k	Born probability	$\text{tr}(\rho E_k)$

Symbol	Name	Interpretation
$S(\rho)$	von Neumann entropy	$-\text{tr}(\rho \log \rho)$
F_{ij}	Fisher information	$\mathbb{E}[\partial_i \log \mathcal{L} \partial_j \log \mathcal{L}]$
g^{FR}	Fisher-Rao metric	Information geometry metric
g^{RSVP}	RSVP metric	Extended accessibility metric
$D(\sigma_1, \sigma_2)$	Trace distance	$\frac{1}{2} \ \sigma_1 - \sigma_2\ _1$
Δ	Laplace-Beltrami	Geometric Laplacian on (\mathcal{M}, g)
\hat{n}	Photon number	$a^\dagger a$
Q	Mandel parameter	Sub/super-Poissonian character
$\Gamma(m)$	Degeneracy ratio	Fiber volume relative to total
Γ_D	Decoherence rate	Rate of distinguishability loss
K_m	KES memory kernel	m -step non-Markovian kernel
\mathfrak{A}	Total admissibility	$\int_X \mathcal{A}(x, t) dx$

Contents

Preface	iii
Notation and Symbols	v
I Mathematical Prerequisites	1
1 Hilbert Spaces, Operators, and Quantum Optical States	3
1.1 Hilbert Spaces and Fock Space	3
1.2 Quantum Optical Measurements	4
2 Manifolds, Riemannian Geometry, and Statistical Manifolds	5
2.1 Smooth Manifolds and Differential Operators	5
2.2 Statistical Manifolds and the Fisher-Rao Metric	6
2.3 Entropy and Variational Principles	6
3 Bayesian Inference: A Rigorous Foundation	7
3.1 Probability Spaces and Bayes' Theorem	7
3.2 Conjugate Priors and Exponential Families	7
3.3 Variational Inference and the ELBO	8

II	Failures of Traditional Models	9
4	The Born Rule and Its Discontents	11
4.1	The Measurement Problem	11
4.2	QBism and the Epistemic Deflation	12
4.3	The Path Integral and Its Measure Problem	13
5	Prior Specification in Quantum Optical Inference	15
5.1	The Prior Assignment Problem	15
5.2	Maximum Entropy Priors and the UV Divergence . . .	16
6	Coherence, Decoherence, and the Entropy of Observation	17
6.1	The Lindblad Master Equation and Its Limits	17
7	Statistical Pathologies in Quantum Optical Tomography	19
7.1	Boundary Concentration of MLE Estimates	19
III	The RSVP Framework	21
8	Admissibility Geometry: Probability as Emergent Accessibility	23
8.1	Trajectory Space and Representational Compression .	23
8.2	Entropy as Accessibility	24
8.3	Why Bayesian Inference Fails at the Boundary of M .	25
8.4	Minimal Working Example: One-Dimensional RSVP System	25
9	The Relativistic Scalar-Vector Plenum: Field Equa-	

tions	27
9.1 The RSVP Lagrangian	27
9.2 Well-Posedness and the RSVP Prior	28
9.3 Correspondence Limits	29
9.4 UV Finiteness of the Multimode Prior	29
10 RSVP Applied to Quantum Optical Inference	31
10.1 The Velocity Field as Natural Gradient	31
10.2 Mandel Q Inference	32
11 RSVP Photon Propagation and Entropic Phase	33
11.1 RSVP-Modified Propagator	33
11.2 Optics as Admissibility Dynamics	34
IV The Spherepop/KES Calculus	35
12 The Spherepop Operators: An Irreversible Event Calculus	37
12.1 Event Primacy	37
12.2 The Four Operators	38
13 KES: Kinetic-Event Synthesis	41
13.1 The KES Map	41
13.2 KES Update and No-Collapse Interpretation	42
13.3 KES Entropy and the Second Law	42
13.4 Non-Markovian Memory	43

V	Synthesis and Connections	45
14	Unification: RSVP Fields as KES Selection Geometry	47
14.1	The Unification Theorem	47
15	Comparisons: QBism, Darwinism, Entropic Dynamics, Free Energy	49
15.1	QBism	49
15.2	Quantum Darwinism	50
15.3	Entropic Dynamics	50
15.4	Friston's Free Energy Principle	50
16	Experimental Predictions and Falsifiability	51
VI	Extensions	53
17	Projection Failure and Representational Collapse	55
17.1	Fiber Degeneracy and Observational Aliasing	55
17.2	Boundary Singularities as Measure-Zero Fibers	56
17.3	Projection Curvature and Berry Phase	56
18	Information Geometry Beyond Fisher-Rao	59
18.1	Limitations of Fisher-Rao Geometry	59
18.2	The RSVP Metric	60
18.3	Admissibility Geodesics	60
19	Non-Markovian Dynamics and Recursive Inference	63
19.1	Inference With Historical Retention	63
19.2	Admissibility Continuity Equation	64

20 Photonics as Constraint Propagation	65
20.1 Optical Systems as Admissibility Transformers	65
21 Adversarial Robustness of Photonic Neural Networks: Non-Ideality as Defense	67
21.1 Hardware Non-Ideality as an Admissibility Constraint	67
21.2 Unary Encoding as Admissibility Restriction	69
21.3 Memory-Efficient Truncated Complementary Unary Encoding	70
21.4 Post-Attack Recovery via Sensitivity-Aware Weight Locking	72
21.5 RSVP Interpretation: Non-Ideality as Admissibility Engineering	73
21.6 Experimental Grounding	75
22 Topological Defects in RSVP Fields	77
22.1 Entropy Vortices and Accessibility Horizons	77
23 RSVP Thermodynamics	79
23.1 Temperature as Accessibility Gradient	79
23.2 Entropy Production and the Admissibility Dissipation Functional	80
24 The Geometry of Learning and Semantic Compression	81
24.1 Learning as Accessibility Compression	81
24.2 CLIO: Constraint-Leveraged Inference and Optimization	82
25 Functorial Structure of Spherepop	83
25.1 Spherepop as a Symmetric Monoidal Category	83
25.2 Collapse as a Functor and Topological Obstructions .	84

26 Computational Complexity of RSVP Inference	85
26.1 Scaling of Field Evolution	85
26.2 Memory Cost and Approximate Inference	86
27 RSVP Cosmological Optics	87
27.1 Entropic Redshift	87
27.2 Structure Formation as Accessibility Condensation . .	88
VII Toward a Unified Accessibility Calculus	89
28 Projection, Agency, and Observer Geometry	91
28.1 Observers as Recursive Projection Systems	91
29 The Unified Accessibility Calculus and Open Problems	93
29.1 The Universal Accessibility Functional	93
29.2 Mathematical Open Problems	94
29.3 The Four Core Ideas	94

Part I

Mathematical Prerequisites

Chapter 1

Hilbert Spaces, Operators, and Quantum Optical States

1.1 Hilbert Spaces and Fock Space

A *Hilbert space* \mathcal{H} is a complete inner product space over \mathbb{C} . The bosonic Fock space for quantum optics is $\mathcal{F}(\mathcal{H}_1) = \bigoplus_{n=0}^{\infty} \mathcal{H}_1^{\otimes_s n}$, where $\mathcal{H}_1 = L^2(\mathbb{R}^3) \otimes \mathbb{C}^2$ is the single-photon Hilbert space and \otimes_s denotes the symmetrized tensor product.

Definition 1.1 (Density operator). A *density operator* $\rho: \mathcal{H} \rightarrow \mathcal{H}$ satisfies $\rho = \rho^\dagger$, $\rho \geq 0$, and $\text{tr}(\rho) = 1$. Pure states satisfy $\rho^2 = \rho$.

The creation and annihilation operators satisfy $[a_k, a_{k'}^\dagger] = \delta_{kk'}$. Coherent states $|\alpha\rangle = e^{-|\alpha|^2/2} \sum_n \frac{\alpha^n}{\sqrt{n!}} |n\rangle$ satisfy $a|\alpha\rangle = \alpha|\alpha\rangle$ and are overcomplete. The Wigner function $W(\alpha) = \frac{2}{\pi} \text{tr}(\rho D(-\alpha)(-1)^{a^\dagger a} D(\alpha))$ is a quasi-probability distribution whose negativity witnesses non-classicality.

1.2 Quantum Optical Measurements

Definition 1.2 (POVM). A *positive operator-valued measure* is a set $\{E_k\}$ with $E_k \geq 0$ and $\sum_k E_k = \mathbf{1}$. The probability of outcome k is $p_k = \text{tr}(\rho E_k)$.

POVMs generalize projection-valued measures to realistic detectors that never project onto sharp eigenstates. Informationally complete POVMs with $M \geq d^2$ elements are required for quantum state tomography of a d -dimensional system.

Chapter 2

Manifolds, Riemannian Geometry, and Statistical Manifolds

2.1 Smooth Manifolds and Differential Operators

An n -dimensional smooth manifold \mathcal{M} is locally homeomorphic to \mathbb{R}^n with smooth transition maps. A Riemannian metric g assigns an inner product to each tangent space $T_p\mathcal{M}$. The key operators used throughout are: gradient ∇f , divergence $\operatorname{div}(X) = \frac{1}{\sqrt{\det g}} \partial_\mu (\sqrt{\det g} X^\mu)$, and Laplace-Beltrami $\Delta f = \operatorname{div}(\nabla f)$.

2.2 Statistical Manifolds and the Fisher-Rao Metric

A *statistical manifold* parameterizes a family of probability distributions $\{p(\cdot | \theta)\}_{\theta \in \mathcal{M}}$. The Fisher-Rao metric is

$$g_{\mu\nu}^{\text{FR}}(\theta) = \int p(x | \theta) \frac{\partial \log p}{\partial \theta^\mu} \frac{\partial \log p}{\partial \theta^\nu} dx. \quad (2.1)$$

Theorem 2.1 (Chentsov, 1972). *Up to a constant, the Fisher-Rao metric is the unique Riemannian metric on the manifold of probability distributions invariant under Markov morphisms.*

The RSVP metric extends the Fisher-Rao metric by incorporating entropy gradients (Chapter 17).

2.3 Entropy and Variational Principles

The Jaynes maximum entropy principle selects the distribution p^* maximizing $H(p) = -\int p \log p$ subject to moment constraints $\mathbb{E}_p[f_k] = c_k$. The solution is always an exponential family $p^*(x) = Z^{-1} e^{-\sum_k \lambda_k f_k(x)}$. The RSVP framework implements this principle *dynamically* via the field equation for S rather than as a static optimization.

Chapter 3

Bayesian Inference: A Rigorous Foundation

3.1 Probability Spaces and Bayes' Theorem

A probability space (Ω, \mathcal{F}, P) consists of a sample space, σ -algebra, and countably additive measure with $P(\Omega) = 1$. Bayes' theorem in measure-theoretic form is

$$\pi(\theta | x) = \frac{p(x | \theta) \pi(\theta)}{\int p(x | \theta') \pi(\theta') d\nu(\theta')}, \quad (3.1)$$

valid when the posterior is proper ($\int \pi(\theta | x) d\nu = 1$) and $p(x) > 0$.

3.2 Conjugate Priors and Exponential Families

A prior $\pi(\theta)$ is *conjugate* to $p(x | \theta)$ if the posterior lies in the same parametric family. For the Poisson likelihood relevant to photon counting, the conjugate prior is $\Gamma(a, b)$ and the posterior given counts $\{n_i\}$ is $\Gamma(a + \sum n_i, b + N)$, with posterior mean $(a + \sum n_i)/(b + N)$.

3.3 Variational Inference and the ELBO

When the posterior is intractable, variational inference maximizes the evidence lower bound

$$\text{ELBO}(\phi) = \mathbb{E}_{q_\phi}[\log p(x | \theta)] - \text{KL}(q_\phi || \pi). \quad (3.2)$$

The ELBO is tight only when $q_\phi = \pi(\cdot | x)$. The RSVP framework avoids the structural assumptions of variational families by treating the prior as a dynamically evolving field rather than a fixed parametric form.

Part II

Failures of Traditional Models

Chapter 4

The Born Rule and Its Discontents

Chapter Dependencies

This chapter uses Definition 1.1 (density operator), Definition 1.2 (POVM), and Hilbert space theory from Chapter 1.

4.1 The Measurement Problem

The Born rule $p_k = \text{tr}(\rho E_k)$ is a postulate, not a theorem. The Schrödinger equation is linear, deterministic, and unitary; measurement outcomes are discrete, stochastic, and irreversible. No continuous deformation of the unitary dynamics produces the Born rule.

Theorem 4.1 (Gleason, 1957). *In a separable Hilbert space of dimension ≥ 3 , every frame function $\mu : \mathcal{P}(\mathcal{H}) \rightarrow [0, 1]$ with $\mu(\mathcal{H}) = 1$ arises from a unique density operator ρ via $\mu(V) = \text{tr}(\rho P_V)$.*

Gleason's theorem shows the Born rule is the *only* consistent assignment; it does not explain why measurement outcomes should

form a frame function. The theorem also fails in dimension two, precisely where single-qubit polarization states live.

Orthodox vs. RSVP

Orthodox assumption: The Born rule $p_k = \text{tr}(\rho E_k)$ is a fundamental postulate appended to the unitary dynamics.

RSVP replacement: Born statistics emerge as a limiting case of the KES selection functional (Theorem 13.3) when RSVP fields are trivial ($S = \text{const}$, $\Phi = \log p$). In non-trivial backgrounds, the effective probability acquires field-dependent corrections.

Orthodox pathology: No dynamical explanation for measurement irreversibility.

RSVP resolution: Irreversibility is primitive in the Spheropop event calculus; the Schrödinger equation is a limit, not the foundation.

4.2 QBism and the Epistemic Deflation

QBism treats ρ as an agent's personal probability assignment, and the Born rule as a consistency norm analogous to Dutch book coherence. The SIC-POVM reformulation gives

$$p(E_k) = (d + 1) \sum_j q_j r(E_k | \pi_j) - 1, \quad (4.1)$$

departing from the classical law of total probability by an additive correction. This correction has no dynamical interpretation in QBism. The RSVP framework provides a field-theoretic account: the correction term is the difference between RSVP-weighted and classical probability sums arising from non-trivial S geometry.

4.3 The Path Integral and Its Measure Problem

The Feynman path integral $\mathcal{A} = \int \mathcal{D}[\gamma] e^{iS[\gamma]/\hbar}$ is the closest existing approach to admissibility geometry. The RSVP framework replaces the oscillatory weight $e^{iS/\hbar}$ with the real, positive admissibility weight $\mathcal{A}(x, t) \propto S e^\Phi$, avoiding the phase-cancellation problem of the Feynman integral while retaining the sum-over-histories structure.

Worked Example

Single-photon polarimetry. A photon with state $\rho = 1/2$ (maximally mixed) measured by a polarizing beam splitter gives $p_H = p_V = 1/2$ by Born. In the RSVP framework, this equal probability arises from a flat entropy field $S = \text{const}$ on the Bloch sphere. When S is non-uniform due to prior optical interactions, the effective probabilities shift by an $O(\alpha S)$ correction testable via photon statistics.

Chapter 5

Prior Specification in Quantum Optical Inference

Chapter Dependencies

This chapter depends on Definition 3.1 (Jeffreys prior), Bayes' theorem (3.1), and Section 3.2 (conjugate priors).

5.1 The Prior Assignment Problem

No canonical rule exists for specifying $\pi(\theta)$ prior to data. The Jeffreys prior $\pi_J(\theta) \propto \sqrt{\det F(\theta)}$ is reparameterization-invariant but concentrates mass near the boundary of the density operator manifold for mixed-state tomography, and is computationally intractable for $n \geq 3$ qubits.

Proposition 5.1 (Positivity violation under Gaussian priors). *Let $\tilde{\rho}$ be drawn from a Gaussian prior on the entries of a $d \times d$ Hermitian matrix with variance $\sigma^2 > 0$. Then $\Pr(\tilde{\rho} \geq 0) \rightarrow 0$ as $d \rightarrow \infty$.*

Proof. The eigenvalues of a GUE matrix with variance σ^2 follow the Wigner semicircle law on $[-2\sigma, 2\sigma]$, which assigns positive mass to

$(-2\sigma, 0)$. By Soshnikov's edge universality (1999), λ_{\min} has Tracy-Widom fluctuations around -2σ , so $\Pr(\lambda_{\min} \geq 0) \rightarrow 0$ as $d \rightarrow \infty$. \square

Orthodox vs. RSVP

Orthodox assumption: Gaussian or Jeffreys priors are standard choices for quantum state tomography, enforced by Cholesky reparameterization.

RSVP replacement: The RSVP prior $\pi_{\text{RSVP}}(\theta) = S(\theta)e^{\Phi(\theta)}/Z$ is strictly positive by Theorem 9.4, derived from field dynamics rather than algebraic constraint.

Orthodox pathology: Positivity must be enforced post hoc, introducing Jacobian biases toward pure states.

RSVP resolution: Positivity is a dynamical consequence of $S > 0$.

5.2 Maximum Entropy Priors and the UV Divergence

The Jaynes maximum entropy multimode prior $Z = \prod_k (1 - e^{-\beta_k})^{-1}$ diverges unless an ultraviolet cutoff is imposed. The RSVP entropy field resolves this via Theorem 9.9: the coupling structure of the field equations provides intrinsic UV regularization.

Chapter 6

Coherence, Decoherence, and the Entropy of Observation

Chapter Dependencies

This chapter depends on the Lindblad equation (6.1), the decoherence functional (Definition 6.1), and Section 2.1 (manifolds).

6.1 The Lindblad Master Equation and Its Limits

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[H, \rho] + \sum_k \gamma_k \left(L_k \rho L_k^\dagger - \frac{1}{2} \{L_k^\dagger L_k, \rho\} \right). \quad (6.1)$$

The pathology is not in the Lindblad form but in the post hoc assignment of L_k and γ_k : these are fit to data, converting a dynamical equation into a statistical interpolation device with no principled connection to the Bayesian update on ρ .

Definition 6.1 (Decoherence functional). For bipartite system-environment evolution U_t from $|\psi_0\rangle_S \otimes |\varepsilon_0\rangle_E$, the decoherence func-

tional for histories α, α' is $D(\alpha, \alpha') = \text{tr}(C_\alpha \rho_0 C_{\alpha'}^\dagger)$ where $C_\alpha = P_{\alpha_n}^{t_n} U_{t_n - t_{n-1}} \cdots P_{\alpha_1}^{t_1} U_{t_1}$. Histories are *consistent* if $D(\alpha, \alpha') \approx 0$ for $\alpha \neq \alpha'$.

The pointer basis problem: the preferred basis for consistency is not selected by the formalism. The Spherpap **Refuse** operator resolves this (Theorem 12.7).

Theorem 6.2 (Entropy non-monotonicity under non-Markovian evolution). *Let $\{\mathcal{E}_t\}$ be a non-divisible family of CPTP maps. Then there exist initial states ρ_1, ρ_2 such that $D(\mathcal{E}_t(\rho_1), \mathcal{E}_t(\rho_2))$ is non-monotone in t .*

Proof. Non-divisibility means $\mathcal{E}_{t,s} = \mathcal{E}_t \circ \mathcal{E}_s^{-1}$ is not completely positive for some $t > s$. By the Choi-Jamiołkowski isomorphism, there exists a bipartite state τ such that $(\mathcal{E}_{t,s} \otimes \text{id})(\tau)$ has a negative eigenvalue. Constructing ρ_1, ρ_2 from its eigenvectors produces a backflow of distinguishability, violating monotone decay. \square

Chapter 7

Statistical Pathologies in Quantum Optical Tomography

Chapter Dependencies

This chapter depends on the MLE objective (7.1), Proposition 5.1, and the density operator manifold structure from Chapter 1.

7.1 Boundary Concentration of MLE Estimates

$$\hat{\rho}_{\text{MLE}} = \arg \max_{\rho \geq 0, \text{tr } \rho = 1} \sum_k n_k \log \text{tr}(\rho E_k). \quad (7.1)$$

Proposition 7.1 (Boundary concentration). *Let $\text{rank}(\rho^*) = r < d$. For finite N , $\hat{\rho}_{\text{MLE}}$ is almost surely rank- d , and Fisher-information confidence regions at $\hat{\rho}_{\text{MLE}}$ are asymptotically incorrect because ρ^* lies on the boundary of the density matrix manifold.*

Proof. The log-likelihood is strictly concave on the interior of the

density matrix body, with gradient diverging to $-\infty$ toward the boundary in any direction with $n_k > 0$. Hence any interior maximum has $\text{rank}(\hat{\rho}_{\text{MLE}}) = d$. Asymptotic normality of MLE requires the true parameter to be interior; boundary targets invalidate the delta method. \square

Orthodox vs. RSVP

Orthodox assumption: Regularization (hedged MLE, MAP with Bures prior) enforces positivity algebraically.

RSVP replacement: Positivity of the posterior emerges from $S > 0$ without regularization (Theorem 9.4).

Orthodox pathology: Regularization Jacobians bias estimates toward pure states.

RSVP resolution: The RSVP prior generates its own curvature through the entropy field, with no Cholesky-induced bias.

Part III

The RSVP Framework

Chapter 8

Admissibility Geometry: Probability as Emergent Accessibility

Chapter Dependencies

This chapter introduces the central conceptual foundation of the RSVP framework. All subsequent chapters in Parts II–VI depend on Definitions 8.1–8.2 and Axiom 8.3.

8.1 Trajectory Space and Representational Compression

Let X be the space of all admissible world-histories and M a finite-dimensional manifold on which statistical models operate. A *projection* $\pi : X \rightarrow M$ maps trajectory structure into observable quantities. For each $m \in M$, the *fiber* $\pi^{-1}(m)$ is the set of all trajectories projecting to the same point. Classical models are defined entirely on M and are silent about fiber structure.

Definition 8.1 (Admissibility functional). An *admissibility functional* $\mathcal{A} : X \times \mathbb{R} \rightarrow [0, \infty)$ measures the degree to which trajectory x remains dynamically and inferentially available at time t .

Definition 8.2 (Emergent probability). The *emergent probability* of a representational event $B \subset M$ is

$$P(B, t) = \frac{\int_{\pi^{-1}(B)} \mathcal{A}(x, t) dx}{\int_X \mathcal{A}(x', t) dx'}. \quad (8.1)$$

Axiom 8.3 (Admissibility primacy). The admissibility functional \mathcal{A} is the fundamental object of inference. Probability distributions over M are derived quantities obtained by projecting \mathcal{A} via equation (8.1).

8.2 Entropy as Accessibility

In the admissibility framework, entropy is the logarithm of accessible trajectory volume:

$$S(x, t) = \log \Omega(x, t), \quad (8.2)$$

where $\Omega(x, t)$ is the volume of admissible continuations from x at time t . Low entropy means strong constraint; high entropy means rich accessibility. This differs from Shannon uncertainty, which is defined over predefined symbol sets.

Proposition 8.4 (Shannon entropy as a limit). *If $\mathcal{A}(x, t) = p(x)$ is time-independent and fibers are uniform, then the accessibility entropy S equals the Shannon entropy of the projected distribution $P(\cdot, t)$ up to a constant depending on fiber volume.*

Proof. Under uniform fibers $|\pi^{-1}(m)| = C$, $\Omega(m, t) = Cp(m)$, so $S(m, t) = \log p(m) + \log C$. Then $-\mathbb{E}[S] = H(P(\cdot, t)) - \log C$, giving $H = \mathbb{E}[S] + \text{const}$. \square

8.3 Why Bayesian Inference Fails at the Boundary of M

The projection π is many-to-one. Inference fails at the boundary of M precisely when the boundary corresponds to measure-zero trajectory fibers: the density matrix manifold boundary (pure states) is the image of a measure-zero set of fibers in trajectory space, yet maximum likelihood and Lindblad machineries assign finite probability to these boundary regions by treating them as interior points of M . This is the geometric explanation for Proposition 7.1.

Orthodox vs. RSVP

Orthodox assumption: The representational manifold M is a complete description; all relevant information is in $m \in M$.

RSVP replacement: M is a projection of X ; the fiber $\pi^{-1}(m)$ contains the lost causal history, alternative admissible paths, and trajectory distinctions relevant to future inference.

Orthodox pathology: Fiber degeneracy causes observational aliasing, boundary singularities, and representational fragility under model change.

RSVP resolution: The admissibility functional is defined on X , making fiber structure explicit and the cost of compression computable via the degeneracy ratio $\Gamma(m)$.

8.4 Minimal Working Example: One-Dimensional RSVP System

Before introducing the full RSVP machinery, we work through a minimal complete example to anchor the formalism.

Worked Example

Binary Bayesian inference via 1D RSVP. Let $\Theta = \{0, 1\}$ encoded as $\mathcal{M} = [-1, 1]$, with $\Phi(\theta) = \log p(D | \theta)$ initialized from a binomial likelihood, $S(\theta) = (1 - \theta^2)^{1/2}$ (high entropy near $\theta = 0$, low near the boundaries), and $\mathbf{v}(\theta) = (\alpha/S)\nabla\Phi$. The RSVP prior is $\pi_{\text{RSVP}}(\theta) \propto (1 - \theta^2)^{1/2}e^{\Phi(\theta)}$, which is a semicircle-weighted likelihood. The posterior concentrates near the maximum of $\Phi(\theta)$ but is moderated by the entropy field: if data strongly favor $\theta = +1$ (a boundary), the entropy field $S(\pm 1) = 0$ suppresses the prior at the boundary, preventing the overconfident boundary concentration of the MLE. A three-event KES history $H_3 = (e_1, e_2, e_3)$ with e_i drawn from this system gives an empirical frequency that converges to π_{RSVP} by Theorem 13.3.

Chapter 9

The Relativistic Scalar-Vector Plenum: Field Equations

Chapter Dependencies

This chapter depends on Chapter 8 (admissibility geometry), Chapter 2 (Riemannian manifolds and Fisher-Rao metric), and the Euler-Lagrange formalism.

9.1 The RSVP Lagrangian

Definition 9.1 (RSVP Lagrangian density).

$$\mathcal{L}_{\text{RSVP}} = \frac{1}{2} |\nabla \Phi|^2 - \frac{1}{2} S g(\mathbf{v}, \mathbf{v}) - V(\Phi, S) + \alpha \Phi \operatorname{div}(\mathbf{v}) - \beta S \log S. \quad (9.1)$$

The term $-\beta S \log S$ drives S dynamically toward maximum entropy configurations, implementing the Jaynes principle as field evolution. The Euler-Lagrange equations are:

$$\Delta \Phi = \frac{\partial V}{\partial \Phi} - \alpha \operatorname{div}(\mathbf{v}), \quad (9.2)$$

$$S \mathbf{v}^\flat = \alpha d\Phi, \quad (9.3)$$

$$\partial_t S = -\operatorname{div}(S\mathbf{v}) + \beta(1 + \log S) + \frac{\partial V}{\partial S}. \quad (9.4)$$

Technical Note

Numerical evolution. Equations (9.2)–(9.4) are discretized on a rectangular lattice with semi-Lagrangian advection for $\operatorname{div}(S\mathbf{v})$. Clamp $S \geq \epsilon > 0$ at each step.

```
Phi += dt*(laplacian(Phi) - dV_dPhi + alpha*divergence(v))
v    = (alpha/S)*gradient(Phi)
S    += dt*(-divergence(S*v) + beta*(1+log(S)) + dV_dS)
S    = clip(S, epsilon, None)
```

9.2 Well-Posedness and the RSVP Prior

Theorem 9.2 (Well-posedness). *Let (\mathcal{M}, g) be compact without boundary, $\dim \leq 4$, and $V = \frac{\lambda}{4}(\Phi^2 - \mu^2)^2 + \frac{\kappa}{2}\Phi S$. For initial data $(\Phi_0, \mathbf{v}_0, S_0) \in H^2 \times H^1 \times H^1$ with $S_0 \geq s_0 > 0$, there exists a unique strong solution on $[0, T]$ with $S \geq s_0 e^{-Ct}$.*

Proof. Define energy $\mathcal{E} = \frac{1}{2}\|\nabla\Phi\|^2 + \frac{1}{2}\|S^{1/2}\mathbf{v}\|^2 + \int V \, d\operatorname{vol} + \beta \int S \log S \, d\operatorname{vol}$. Differentiating and integrating by parts on the compact manifold eliminates boundary terms. Sobolev embedding $H^2 \hookrightarrow L^\infty$ for $n \leq 4$ bounds the potential gradient. Gronwall gives $\mathcal{E}(t) \leq \mathcal{E}(0)e^{C_1 t}$; local existence follows from contraction mapping. The S lower bound follows from the comparison principle applied to (9.4). \square

Definition 9.3 (RSVP-generated prior). At steady state, the RSVP prior over $\Theta \subset \mathcal{M}$ is

$$\pi_{\text{RSVP}}(\theta) = \frac{S(\theta) e^{\Phi(\theta)}}{\int_{\Theta} S(\theta') e^{\Phi(\theta')} \, d\operatorname{vol}_g}. \quad (9.5)$$

Theorem 9.4 (RSVP prior positivity). *Under Theorem 9.2, π_{RSVP} is strictly positive on Θ and defines a probability measure, resolving Proposition 5.1 at the dynamical level.*

Proof. $S(\theta)e^{\Phi(\theta)} > 0$ everywhere by $S > 0$ and $e^{\Phi} > 0$. The denominator is a positive continuous integral over a compact set, hence finite and positive. \square

9.3 Correspondence Limits

Proposition 9.5 (Classical Bayesian limit). *If $S = \text{const}$ and $\Phi = \log \mathcal{L}(\theta)$, then RSVP inference reduces to ordinary Bayesian updating with likelihood \mathcal{L} and uniform prior.*

Proposition 9.6 (Shannon entropy limit). *If $\mathcal{A}(x, t) = p(x)$ (time-independent, uniform fibers), then the RSVP entropy field S equals the Shannon entropy of the projected distribution up to a fiber-volume constant (Proposition 8.4).*

Proposition 9.7 (Markovian limit). *The KES memory kernel K_m (Definition 13.6) reduces to the standard KES map when $m = 0$, giving Markovian event selection.*

Proposition 9.8 (Gaussian approximation limit). *Near a maximum θ^* of Φ , Taylor expanding to second order gives $\pi_{\text{RSVP}}(\theta) \approx S(\theta^*)e^{\Phi(\theta^*)} \exp(-\frac{1}{2}(\theta - \theta^*)^T F(\theta^*)(\theta - \theta^*))$, recovering a Gaussian posterior with Fisher information covariance.*

9.4 UV Finiteness of the Multimode Prior

Theorem 9.9 (UV finiteness). *Let $\mathcal{M} = \prod_{k=1}^K [0, \infty)$ parameterize K modes. With $\kappa_k \geq ck^\gamma$ for $c, \gamma > 0$, the steady-state RSVP entropy field satisfies $\int_{\mathcal{M}} S^* d\bar{n} < \infty$ without any ultraviolet cutoff.*

Proof. The potential penalizes large \bar{n}_k with strength κ_k , forcing $S^* \leq C \prod_k e^{-\kappa_k \bar{n}_k}$ by the maximum principle. Then $\int S^* \leq C \prod_k \kappa_k^{-1}$, which vanishes as $K \rightarrow \infty$ for $\kappa_k \geq ck^\gamma$. \square

*CHAPTER 9. THE RELATIVISTIC SCALAR-VECTOR PLENUM:
FIELD EQUATIONS*

Chapter 10

RSVP Applied to Quantum Optical Inference

Chapter Dependencies

This chapter applies Theorems 9.2–9.9 to concrete photonic inference problems. Requires the Mandel Q parameter and photon number operator from Chapter 1.

10.1 The Velocity Field as Natural Gradient

Theorem 10.1 (RSVP velocity as entropic natural gradient). *At steady state with $g = g^{\text{FR}}$, the RSVP velocity field satisfies $\mathbf{v} = (\alpha/S)(g^{\text{FR}})^{-1}\nabla\Phi$: the Fisher-Rao natural gradient of the likelihood, suppressed by entropy density in high-uncertainty regions. Bayesian updating is geometrically implemented by the RSVP velocity field. The RSVP framework does not add Bayesian inference to a field theory; it shows that Bayesian inference is a field theory with Lagrangian (9.1).*

Proof. From (9.3), $S\mathbf{v}^b = \alpha d\Phi$. Raising with $(g^{\text{FR}})^{-1}$ gives $S\mathbf{v} = \alpha\nabla_{g^{\text{FR}}}\Phi$, so $\mathbf{v} = (\alpha/S)(g^{\text{FR}})^{-1}\nabla\Phi$. \square

10.2 Mandel Q Inference

Worked Example

Sub-Poissonian light. From $N = 20$ counts with $\bar{n} = 0.4$ and $s^2 = 0.1$, the standard estimator gives $\hat{Q} = -0.75$. A Gaussian prior on $Q \in (-1, \infty)$ assigns non-negligible mass to the unphysical region $Q < -1$. The RSVP prior with $V \supset -\lambda \log(Q+1)$ gives $\Phi \rightarrow -\infty$ as $Q \rightarrow -1^+$, automatically enforcing the constraint. The posterior mean is regularized toward $Q = 0$ by S weighting, preventing over-interpretation of small-sample fluctuations.

Chapter 11

RSVP Photon Propagation and Entropic Phase

Chapter Dependencies

This chapter depends on Theorem 9.4, Section 2.1 (manifolds), and the free photon propagator from quantum field theory.

11.1 RSVP-Modified Propagator

Definition 11.1 (RSVP-modified propagator).

$$\tilde{D}_{\mu\nu}(x, y) = D_{\mu\nu}(x - y) \exp\left(-\frac{1}{2} \int_x^y \nabla\Phi \cdot dl - \frac{\alpha}{2} \int_x^y S ds\right). \quad (11.1)$$

The potential factor $e^{-(\Phi(y)-\Phi(x))/2}$ is path-independent. The attenuation factor $e^{-\frac{\alpha}{2} \int S ds}$ encodes decoherence without invoking the environment explicitly: photons traversing high-entropy regions experience additional damping.

Proposition 11.2 (RSVP entropic phase shift). *A photon of frequency ω_0 propagating from x_0 to x_1 acquires $\Delta\omega = \frac{\omega_0}{2}(\Phi(x_1) - \Phi(x_0)) + O(\alpha S)$.*

Epistemic Status

The identification of $\tilde{D}_{\mu\nu}$ with the physical photon Green's function is an ansatz. Testing the entropic phase shift requires a regime where $\nabla\Phi$ is significant along the photon path, e.g., photonic crystals with engineered entropy profiles.

11.2 Optics as Admissibility Dynamics

Standard optical phenomena acquire RSVP interpretations: interference is admissibility-bundle compatibility; diffraction is the spread of admissible continuations after an aperture; coherence $g^{(1)}(\tau)$ measures admissibility-bundle overlap over time τ ; a lens is an admissibility operator imposing $\Phi(x) = -k|x|^2/(2f)$ that drives \mathbf{v} toward focal convergence.

Worked Example

Mach-Zehnder interferometer. A photon entering path 1 or path 2 corresponds to two trajectory bundles x_1, x_2 with $\pi(x_1) = \pi(x_2)$ at the final detector (same outcome). Constructive interference occurs when the admissibility phases are compatible, i.e., $\Phi(x_1) - \Phi(x_2) \in 2\pi\mathbb{Z}$. Destructive interference corresponds to admissibility cancellation: the fiber $\pi^{-1}(\text{detector})$ has two contributions that exactly cancel in the admissibility integral (8.1).

Part IV

The Spheredop/KES Calculus

Chapter 12

The SpheroPOP Operators: An Irreversible Event Calculus

Chapter Dependencies

This chapter depends on Chapter 8 (admissibility geometry) and the decoherence functional (Definition 6.1).

12.1 Event Primacy

Axiom 12.1 (Event primacy). A state in the usual sense is an equivalence class of event histories. Two histories are equivalent if and only if they are indistinguishable by all subsequent events.

Axiom 12.1 inverts the standard ordering: states are derived from events, not the reverse. Irreversibility is primitive, not emergent.

Definition 12.2 (Event algebra). An *event algebra* $(\mathcal{E}, \leq, \sqcup, \perp)$ consists of events, causal precedence, event composition, and null event.

An event e is *irreversible* if no e^{-1} satisfies $e \sqcup e^{-1} = \perp$. All photon detection events are irreversible.

12.2 The Four Operators

Definition 12.3 (Pop). Maps an event to its immediate causal consequence; models photon detection as the creation of a new event, not a state update.

Definition 12.4 (Bind). Joins causally compatible events: $\mathbf{Bind}(e_1, e_2) = e_1 \sqcup e_2$ when neither $e_1 \leq e_2$ nor $e_2 \leq e_1$; models coincidence measurement of entangled pairs.

Definition 12.5 (Collapse). Maps an event e to the subspace $\{|\psi\rangle : P_k|\psi\rangle \neq 0\}$ of states compatible with outcome k ; identifies compatible configurations without destroying pre-event information.

Definition 12.6 (Refuse). Filters events by a consistency condition \mathcal{C} : $\mathbf{Refuse}(A, \mathcal{C}) = \{e \in A : e \text{ is consistent with all } c \in \mathcal{C}\}$; implements pointer basis selection.

Theorem 12.7 (Refuse resolves the pointer basis problem). *With $\mathcal{C} = \{\text{robustness under tracing out } \mathcal{H}_E\}$, $\mathbf{Refuse}(\mathcal{E}_S, \mathcal{C})$ selects precisely the Quantum Darwinism pointer basis: system states that leave redundant records in multiple environmental fragments.*

Proof. A state is in the pointer basis iff $I(S : E_\alpha) \approx H(S)$ for many disjoint environmental fragments \mathcal{H}_{E_α} . The **Refuse** consistency condition requires any event associated with the state to be reproducible from multiple fragments, which is precisely the Quantum Darwinism redundancy condition. \square

Proposition 12.8 (Non-commutativity of Pop and Bind). $\mathbf{Pop}(\mathbf{Bind}(e_1, e_2)) \neq \mathbf{Bind}(\mathbf{Pop}(e_1), \mathbf{Pop}(e_2))$ in general.

Proof. For photons from parametric down-conversion, $\mathbf{Pop}(\mathbf{Bind}(e_1, e_2))$ yields HOM-correlated outcomes; $\mathbf{Bind}(\mathbf{Pop}(e_1), \mathbf{Pop}(e_2))$ yields no

HOM interference. The measurement statistics differ experimentally, confirming the inequality. \square

*CHAPTER 12. THE SPHEREPOP OPERATORS: AN IRREVERSIBLE
EVENT CALCULUS*

Chapter 13

KES: Kinetic-Event Synthesis

Chapter Dependencies

This chapter depends on the Spherepop operators (Chapter 12), the RSVP prior (9.5), and the admissibility functional (Definition 8.1).

13.1 The KES Map

Definition 13.1 (World-state and history state). $\Omega_t = (\mathcal{X}_t, w_t)$ with $w_t : \mathcal{X}_t \rightarrow [0, 1]$ normalized is the posterior-weighted set of consistent configurations at time t . $H_t = (e_1, \dots, e_m)$ is an irreversible causal sequence.

Definition 13.2 (KES map). The KES map $\text{KES} : \Omega_t \rightarrow H_{t+1}$ selects e^* maximizing

$$\mathcal{F}(e) = w_t(\mathbf{Collapse}(e)) \cdot S(\mathbf{Collapse}(e)) \cdot e^{\Phi(\mathbf{Collapse}(e))}, \quad (13.1)$$

and sets $H_{t+1} = H_t \sqcup (e^*)$.

Theorem 13.3 (KES recovers Born statistics). *When $S = \text{const}$ and $\Phi = \log p$, the long-run frequency of KES-selected outcomes converges almost surely to $p(e_k) = \text{tr}(\rho E_k)$.*

Proof. $\mathcal{F}(e_k) = S_0 w_t(\mathbf{Collapse}(e_k)) p(e_k)$; for uniform w_t , KES samples from $(p(e_k))_k$. The strong law of large numbers gives almost-sure convergence to $\text{tr}(\rho E_k)$. \square

13.2 KES Update and No-Collapse Interpretation

Post-selection updates w_t by Bayesian conditioning on e^* :

$$w_{t+1}(x) = \frac{w_t(x) \cdot \mathbf{1}[x \in \mathbf{Collapse}(e^*)]}{\sum_{x' \in \mathbf{Collapse}(e^*)} w_t(x')}. \quad (13.2)$$

No Hilbert-space collapse occurs. The apparent collapse is an artifact of conflating history states with world-states.

13.3 KES Entropy and the Second Law

Definition 13.4 (KES entropy).

$$\mathcal{S}_{\text{KES}}(t) = - \sum_x w_t(x) \log w_t(x) + |H_t| \log |\mathcal{E}|. \quad (13.3)$$

Theorem 13.5 (KES second law). $\mathbb{E}[\mathcal{S}_{\text{KES}}(t+1)] \geq \mathcal{S}_{\text{KES}}(t)$, with equality iff w_t is a point mass.

Proof. The history term increases by $\log |\mathcal{E}| > 0$ per step. The Shannon term decreases by at most $H(w_t)$ by the data processing inequality. The net change is non-negative for $|\mathcal{E}| \geq |\mathcal{X}_t|$. \square

13.4 Non-Markovian Memory

Definition 13.6 (KES memory kernel of depth m).

$$K_m(e_{t-m}, \dots, e_{t-1}) = \int_{\mathcal{X}_t} w_t(x | H_{t-m:t}) \mathcal{F}(e^*(x)) dx. \quad (13.4)$$

For $m = 0$, the kernel is Markovian. For $m > 0$, past events influence event selection, modeling the distinguishability revivals of Theorem 6.2.

Technical Note

KES memory as transformer attention. The memory kernel K_m is approximated by attention over the last m events encoded as feature vectors. This maps KES dynamics directly onto transformer architecture: attention heads may be learning approximations to admissibility kernels.

Part V

Synthesis and Connections

Chapter 14

Unification: RSVP Fields as KES Selection Geometry

Chapter Dependencies

This chapter synthesizes all preceding material. Requires Theorems 9.2, 9.4, 13.3, and the Spherepop operators.

14.1 The Unification Theorem

Theorem 14.1 (RSVP-KES unification). *Let $(\mathcal{M}, g, (\Phi, \mathbf{v}, S))$ be a compact RSVP configuration and define the KES functional by (13.1). Then: (1) KES-selected events form a Markov chain with kernel $K(e, e') = \mathcal{F}(e')/Z(e)$; (2) the stationary distribution is π_{RSVP} ; (3) convergence is exponential with rate given by the spectral gap of the RSVP Laplacian.*

Proof. (1) $\mathcal{F}(e') > 0$ and $Z(e) < \infty$ on a compact manifold give a well-defined kernel. (2) At steady state $\mathcal{F}(\theta) \propto S e^\Phi$ and Z is constant; detailed balance confirms $\pi = \pi_{\text{RSVP}}$. (3) The Poincaré inequality on a compact Riemannian manifold gives a positive spectral gap for Δ ,

which controls the Markov chain mixing time.

□

Chapter 15

Comparisons: QBism, Darwinism, Entropic Dynamics, Free Energy

Chapter Dependencies

This chapter compares RSVP-KES to four existing programs. Requires Theorems 13.3, 12.7, and the ELBO (3.2).

For each framework, we state what RSVP-KES reproduces, what it modifies, and what predictions differ.

15.1 QBism

Reproduced: Epistemic status of ρ ; Born-rule consistency in the flat limit (Theorem 13.3). **Modified:** The consistency norm is replaced by the KES functional dynamics. The SIC-POVM correction term acquires a field-theoretic account via S geometry. **Differing predictions:** In non-trivial RSVP backgrounds, effective probabilities deviate from Born by $O(\alpha S)$. This is absent in QBism, which provides

no field dynamics.

15.2 Quantum Darwinism

Reproduced: Theorem 12.7 shows **Refuse**($\mathcal{E}_S, \mathcal{C}_{\text{robust}}$) selects precisely the pointer basis. **Modified:** The RSVP entropy field S provides a dynamical measure of redundancy; low S corresponds to high environmental recording. **Differing predictions:** A specific functional form for the quantum-to-classical transition as a function of S along the system-environment interaction.

15.3 Entropic Dynamics

Reproduced: The term $-\beta S \log S$ in the RSVP Lagrangian implements maximum entropy dynamics, paralleling Caticha's entropic update. **Modified:** Entropic Dynamics does not couple entropy to Φ or \mathbf{v} . The RSVP velocity field has no analog in Entropic Dynamics. **New structure:** The KES memory kernel (13.4) handles non-Markovian dynamics that Entropic Dynamics cannot model.

15.4 Friston's Free Energy Principle

Reproduced: ELBO maximization (3.2) corresponds to minimizing $\text{KL}(q_\phi \parallel \pi_{\text{RSVP}})$. Active inference corresponds to KES event selection. **Modified:** The FEP treats the generative model p as fixed. In RSVP-KES, the prior π_{RSVP} is dynamically generated; evidence updates both the posterior and the prior geometry. **Formal connection:** The FEP is the special case of RSVP-KES where the fields are static (no field dynamics, only posterior updates).

Chapter 16

Experimental Predictions and Falsifiability

Chapter Dependencies

This chapter requires the RSVP propagator (11.1), the KES functional, and the memory kernel (13.4). Predictions 1–3 depend on ansätze stated explicitly.

Prediction 1 *Modified photon detection statistics.* In a photonic system where S is engineered to be non-uniform, KES predicts $\tilde{p}_k \propto S_k \text{tr}(\rho E_k)$, deviating from Born by an $O(S)$ factor. *Ansatz:* identification of \mathcal{F} with physical detection weight.

Prediction 2 *Entropic phase shift.* Photons propagating through a medium with $\nabla\Phi \neq 0$ acquire $\Delta\omega = \frac{\omega_0}{2} \nabla\Phi \cdot \Delta x$, distinguishable from dispersive shifts by its $\nabla\Phi$ dependence.

Prediction 3 *Controlled non-Markovian entropy revivals.* Systems with $m > 0$ KES memory exhibit distinguishability revivals at times determined by K_m without free parameters once S and Φ are calibrated.

Prediction 4 *UV-finite multimode prior.* The RSVP prior is normalizable without UV cutoff (Theorem 9.9), predicting specific mode-number dependence of prior probability testable against Jeffreys and Bures priors.

Prediction 5 *Pointer basis selection by Refuse.* In quantum eraser experiments, **Refuse** predicts which basis is selected with a quantitative relationship between S and the redundancy fraction, agreeing with Quantum Darwinism but providing a specific dynamical mechanism.

Epistemic Status

Predictions 1 and 2 require experimentally accessing RSVP field quantities independently. Current proposals: (a) inverse RSVP field equations from posterior data; (b) deliberate engineering of S via structured reservoir coupling; (c) estimating Φ from the natural gradient structure of the experimental likelihood. These are research directions, not established methods.

Part VI

Extensions

Chapter 17

Projection Failure and Representational Collapse

Chapter Dependencies

This chapter depends on Chapter 8 (admissibility geometry and the projection $\pi : X \rightarrow M$), Proposition 7.1, and Theorem 6.2.

17.1 Fiber Degeneracy and Observational Aliasing

Define the *degeneracy ratio*

$$\Gamma(m) = \frac{\mu(\pi^{-1}(m))}{\int_M \mu(\pi^{-1}(m')) dm'}, \quad (17.1)$$

measuring the relative volume of the fiber over $m \in M$. Large $\Gamma(m)$ means many distinct trajectories are aliased to the same representational point: the inference system cannot distinguish them, and any prediction that depends on their difference is compromised.

Proposition 17.1 (Degeneracy implies inference failure). *If $\Gamma(m) > \Gamma(m')$ and future predictions depend on trajectories in $\pi^{-1}(m) \setminus$*

$\pi^{-1}(m')$, then any model operating on M cannot achieve the prediction accuracy of a model operating on X .

Proof. By definition, the fiber $\pi^{-1}(m)$ contains trajectories that are indistinguishable from m alone. If these trajectories lead to different future outcomes, then any model that uses only m cannot distinguish the outcomes and must average over the fiber, producing prediction error equal to the variance of outcomes within the fiber. \square

17.2 Boundary Singularities as Measure-Zero Fibers

The boundary of M (e.g., pure states in quantum tomography) corresponds to measure-zero fibers in X : only a single trajectory (or a lower-dimensional submanifold) projects to each boundary point. The MLE boundary concentration (Proposition 7.1) is therefore a symptom of a deeper pathology: the inference system is assigning finite probability to representational regions whose preimage in trajectory space has measure zero.

17.3 Projection Curvature and Berry Phase

The curvature of the fiber bundle $\pi : X \rightarrow M$ is measured by the *connection* induced by the RSVP vector field \mathbf{v} on the horizontal distribution $H_x X = \ker(d\pi_x)^\perp$. Non-trivial curvature—the holonomy of this connection around closed loops in M —corresponds to the Berry phase in quantum mechanics: a photon transported around a loop in parameter space acquires a geometric phase proportional to the fiber bundle curvature.

Conjecture 17.2. The Berry phase of a quantum optical system is equal to the holonomy of the RSVP admissibility connection ∇^{RSVP} around the corresponding loop in the statistical manifold M .

Technical Note

Sparse tensor representation of high-dimensional fibers.

For K -mode systems, X has dimension $O(4^K)$ while M has dimension $O(K)$. The fiber $\pi^{-1}(m)$ is an $O(4^K - K)$ -dimensional manifold. Sparse tensor formats (Tucker decomposition, tensor train) can represent the admissibility measure on fibers with $O(Kr^2)$ parameters for bond dimension r , enabling scalable fiber-aware inference.

*CHAPTER 17. PROJECTION FAILURE AND REPRESENTATIONAL
COLLAPSE*

Chapter 18

Information Geometry Beyond Fisher-Rao

Chapter Dependencies

This chapter extends Chapter 2 (Fisher-Rao geometry) using the RSVP entropy field from Chapter 8 and the field equations (9.2)–(9.4).

18.1 Limitations of Fisher-Rao Geometry

The Fisher-Rao metric assumes: fixed probability spaces, predefined observables, and static admissibility. It is a metric on *representational* manifold M , not on the trajectory space X . It captures the geometry of probability distributions but not the geometry of the fibers that generate them.

18.2 The RSVP Metric

Definition 18.1 (RSVP metric). The RSVP metric on \mathcal{M} is

$$g_{ij}^{\text{RSVP}} = g_{ij}^{\text{FR}} + \lambda \partial_i S \partial_j S, \quad (18.1)$$

extending the Fisher-Rao metric by a term encoding entropy gradients.

Entropy gradients contribute to the metric because regions where S changes rapidly are informationally distinct: inference is harder across entropy-gradient boundaries. The RSVP metric penalizes such boundaries geometrically.

Proposition 18.2 (RSVP metric is positive definite). *For $\lambda \geq 0$, g^{RSVP} is positive semidefinite; for $\lambda > 0$ and S non-constant, it is positive definite.*

Proof. g^{FR} is positive semidefinite by construction. The second term $\lambda \partial_i S \partial_j S$ is a rank-one positive semidefinite correction. For $\lambda > 0$ and $\nabla S \neq 0$, the correction is positive definite in the direction of ∇S , and $g^{\text{FR}} > 0$ in the orthogonal directions. \square

18.3 Admissibility Geodesics

Geodesics in the RSVP metric (18.1) generalize Fisher-Rao geodesics (which correspond to the most informatively efficient paths between distributions) by penalizing paths that cross entropy-gradient barriers. These geodesics correspond to:

- Bayesian natural gradient paths (in the Fisher-Rao component),
- Entropy-avoiding inference trajectories (in the $\lambda \partial S \partial S$ component),
- Wasserstein gradient flows when the metric is the L^2 Wasserstein metric on the space of probability measures.

Worked Example

Geodesic on the Bloch sphere. For a single qubit, the Fisher-Rao metric on the Bloch sphere is the round metric $g^{\text{FR}} = d\theta^2 + \sin^2 \theta d\phi^2$. The RSVP metric adds a term $\lambda(\partial_\theta S)^2 d\theta^2 + \lambda(\partial_\phi S)^2 d\phi^2$. If $S(\theta) = \sin \theta$ (maximum entropy at the equator, zero at poles), geodesics are deflected toward the equatorial plane, reflecting the higher admissibility near $\theta = \pi/2$ and the lower admissibility (zero entropy) at the poles (pure states).

Chapter 19

Non-Markovian Dynamics and Recursive Inference

Chapter Dependencies

This chapter depends on Theorem 6.2, the KES memory kernel (Definition 13.6), and Section 2.3 (entropy and variational principles).

19.1 Inference With Historical Retention

The KES memory kernel K_m from Definition 13.6 defines recursive Bayesian accessibility: the admissibility at time t is modulated by the event history $H_{t-m:t}$,

$$P_t(B) = \frac{\int_{\pi^{-1}(B)} \mathcal{A}(x, t, H_t) dx}{\int_X \mathcal{A}(x', t, H_t) dx'}. \quad (19.1)$$

This generalizes equation (8.1) to history-dependent admissibility, enabling systems that fail under Markovian approximation—photonic memory, biological adaptation, recursive learning—to be modeled within the RSVP framework.

19.2 Admissibility Continuity Equation

Theorem 19.1 (Admissibility conservation). *If $\operatorname{div}(\mathbf{v}) = 0$, the total admissibility measure $\mathfrak{A}(t) = \int_{\mathcal{M}} \mathcal{A}(x, t) dx$ is conserved: $\partial_t \mathfrak{A} = 0$.*

Proof. The admissibility continuity equation is $\partial_t \mathcal{A} + \operatorname{div}(\mathcal{A}\mathbf{v}) = 0$. Integrating over \mathcal{M} and applying the divergence theorem gives $\partial_t \mathfrak{A} = -\int_{\partial\mathcal{M}} \mathcal{A}\mathbf{v} \cdot \hat{n} dA$. On a compact manifold without boundary, or with $\mathcal{A}\mathbf{v} \cdot \hat{n} = 0$ on $\partial\mathcal{M}$, the boundary term vanishes. \square

Chapter 20

Photonics as Constraint Propagation

Chapter Dependencies

This chapter applies the RSVP propagator (11.1) and the admissibility framework (Chapter 8) to optical systems from Chapter 1.

20.1 Optical Systems as Admissibility Transformers

An optical system (beam splitter, cavity, lens, fiber) is an operator on the admissibility field that transforms (Φ, \mathbf{v}, S) while preserving certain conserved quantities.

Beam splitters partition the trajectory admissibility: a 50:50 beam splitter performs $\mathcal{A}_{\text{out}} \propto \mathcal{A}_{\text{in}}^{(1)} + \mathcal{A}_{\text{in}}^{(2)}$ with interference structure in the admissibility overlap.

Optical cavities create low-entropy trajectory basins: resonant frequencies correspond to stable fixed points of the RSVP velocity

field \mathbf{v} , where $\mathbf{v} \cdot \nabla\Phi = 0$ (no net flow) and S is locally minimized (high constraint).

Wigner negativity corresponds to inaccessible classical fibers: the negative regions of $W(\alpha)$ are those where classical admissibility \mathcal{A}_{cl} is suppressed but quantum admissibility \mathcal{A}_{qu} is not, reflecting the inaccessibility of the classical trajectory structure in those phase-space regions.

Worked Example

Double-slit decoherence. In a double-slit experiment, two trajectory bundles $x_{\text{slit } 1}$ and $x_{\text{slit } 2}$ arrive at the same screen point. Their admissibility overlap $\int \mathcal{A}(x_{\text{slit } 1}, t) \mathcal{A}(x_{\text{slit } 2}, t) dt$ determines the fringe visibility. As the entropy field S between the slits and the screen increases (due to environmental coupling), the overlap decreases. The decoherence rate is $\Gamma_D = \partial_t D(\rho_1, \rho_2) \propto \alpha \int S ds$, which is the entropy-attenuation term in the RSVP propagator (11.1).

Chapter 21

Adversarial Robustness of Photonic Neural Networks: Non-Ideality as Defense

Chapter Dependencies

This chapter depends on Chapter 8 (admissibility geometry and the projection $\pi : X \rightarrow M$), Chapter 20 (photonics as constraint propagation), and the KES sensitivity functional (13.1). It draws on the work of Lu et al. (2024), which demonstrated that hardware non-idealities in photonic analog neural networks can function as intrinsic adversarial defenders rather than merely as sources of degradation.

21.1 Hardware Non-Ideality as an Admissibility Constraint

Analog optical neural networks (ONNs) are photonic hardware systems in which neural network weights are encoded as light ampli-

tudes using digital-to-analog conversion (DAC) modules, propagated through photonic tensor cores, and accumulated at photodetectors. The standard design concern is that hardware non-idealities—low-precision control, optical noise, crosstalk, fabrication variation—degrade inference accuracy relative to the ideal digital computation. The conventional response is to suppress these non-idealities through hardware-software co-design.

The RSVP framework suggests a different interpretation. Hardware non-idealities are constraints on the admissibility field of the photonic system: they reduce the set of reachable states \mathcal{X}_i and alter the entropy field S by restricting accessible trajectory continuations. From the RSVP perspective, a system with low precision has a coarser admissibility resolution—fewer distinct states are reachable—which is a reduction in S locally. This reduced accessibility is conventionally undesirable for accuracy, but it simultaneously restricts the space available to an adversary performing a weight attack.

Lu et al. (2024) made this precise for a specific class of attacks: gradient-based bit-flip attacks (BFA) on stored neural network weights. Their central observation is that low-bit quantization, sparse (pruned) weight representations, and the unary encoding natively used by optical DAC hardware each reduce the *weight sensitivity* to bit-flip perturbations, and that this reduction has an adversarial protection effect that can be harnessed deliberately. We formalize this observation within the RSVP framework and extend it to a general principle.

Definition 21.1 (Weight sensitivity in the ONN context). Let $\mathcal{L}(W, \mathcal{D})$ be the loss of an ONN with weights W on dataset \mathcal{D} . The *bit-flip sensitivity* of weight W_i is the Taylor expansion of the loss change induced by flipping the most significant bit (MSB) of W_i :

$$S_i = \mathcal{L}(W \oplus \Delta W_i^{\text{MSB}}) - \mathcal{L}(W) \approx \nabla_{W_i} \mathcal{L} \cdot \Delta W_i^{\text{MSB}} + \frac{1}{2} \nabla_{W_i}^2 \mathcal{L} \cdot (\Delta W_i^{\text{MSB}})^2, \quad (21.1)$$

where ΔW_i^{MSB} is the weight perturbation caused by an MSB flip.

In binary-coded decimal (BCD) representation with b bits, an MSB flip changes the weight by approximately half the weight range:

$|\Delta W_i^{\text{MSB}}| \approx 2^{b-1}$. This is a large perturbation, and S_i is correspondingly large for most weights. The adversarial attack selects the weight with the largest $|S_i|$ and flips it, repeating this process until all of a Hamming distance budget HD is exhausted.

21.2 Unary Encoding as Admissibility Restriction

The key hardware observation is that optical DACs in photonic accelerators natively encode weights in *unary* format rather than BCD. In a unary-encoded b -bit weight, the value w is represented as w leading ones followed by $(2^b - 1 - w)$ trailing zeros: $(w)_B = \underbrace{(1, \dots, 1)}_w, \underbrace{(0, \dots, 0)}_{2^b-1-w} U$.

Proposition 21.2 (Unary encoding minimizes MSB sensitivity). *In the unary representation, every bit is a least-significant bit (LSB) in the sense that flipping any single bit changes the weight value by exactly ± 1 (one count unit). In BCD representation, flipping the MSB changes the weight by 2^{b-1} . Therefore, $|\Delta W_i^{\text{MSB}}|_{\text{unary}} \ll |\Delta W_i^{\text{MSB}}|_{\text{BCD}}$ for $b \geq 2$, and the sensitivity S_i is correspondingly reduced by a factor of order 2^{b-1} .*

Proof. In a b -bit BCD representation with 2's complement, the MSB contributes value -2^{b-1} (sign bit). An MSB flip changes the weight by $\pm 2^{b-1}$. In unary representation of a $(2^b - 1)$ -bit string, each bit contributes value $+1$ (one count unit). Any single-bit flip changes the encoded count by exactly ± 1 . Since the weight range is the same in both representations, the ratio of MSB-flip perturbation sizes is $2^{b-1} : 1$. For $b = 8$, this is a $128\times$ reduction in $|\Delta W^{\text{MSB}}|$ and therefore in the first-order sensitivity contribution. \square

In RSVP terms, the unary encoding imposes a constraint on the admissibility field of the weight space: the topology of the weight manifold \mathcal{M}_W changes from one where large jumps are accessible

via single bit-flips (BCD) to one where only unit steps are reachable (unary). This is a reduction in the local admissibility radius of the adversary’s action space. The entropy field S of the weight manifold, which measures the volume of accessible perturbations at each weight, is dramatically lower under unary encoding.

Orthodox vs. RSVP

Orthodox assumption: Hardware non-idealities (low precision, sparse encoding, unary representation) are constraints that degrade accuracy and should be minimized.

RSVP replacement: Non-idealities are admissibility constraints on the weight manifold \mathcal{M}_W . They reduce S locally, restricting the trajectory space available to an adversary. The reduction in adversarial accessibility is the dual of the reduction in representational capacity.

Orthodox pathology: The adversary’s action space (all possible MSB flips) is large in BCD format; small Hamming distance attacks can cause catastrophic accuracy drops.

RSVP resolution: Unary encoding contracts the admissibility radius of the adversary’s actions to $O(1)$ rather than $O(2^{b-1})$, making even large Hamming distance budgets ($HD = 100$) nearly ineffective.

21.3 Memory-Efficient Truncated Complementary Unary Encoding

The original unary representation requires $2^b - 1$ bits to encode a b -bit weight, a factor of $2^{b-1} - 1/b$ memory overhead relative to BCD. This is prohibitive for large networks. Lu et al. (2024) propose a memory-efficient variant, the *truncated complementary unary* (TCU) encoding, which exploits the Gaussian-like weight distribution of trained neural networks.

The key observation is that the trailing zeros of small positive

values and the leading ones of small negative values are redundant for storage: they carry no information beyond alignment. Truncating them reduces the required bitwidth from $2^b - 1$ to \hat{b} , where \hat{b} is chosen to cover the weight’s magnitude. For negative values, the complementary encoding stores trailing zeros (counting from the right) rather than leading ones, achieving the same compression for the negative half of the weight distribution.

Proposition 21.3 (TCU memory overhead). *For a weight W_i with $|W_i| = k$ (in integer count units, $0 \leq k \leq 2^{b-1}$), the TCU representation requires*

$$\hat{b}_i = \lceil \log_2 \min(2^b - |W_i|, |W_i|) \rceil + 1 \quad (21.2)$$

bits, compared to $2^b - 1$ bits for the full unary representation. For weights concentrated near zero (as in trained networks), $\hat{b}_i \ll 2^b - 1$ for most i , giving memory overhead reduction of 6–11× relative to full unary encoding.

The TCU encoding creates an exponentially-spaced partition of the weight range: weights of magnitude k require $O(\log_2 k)$ bits. This is analogous to logarithmic quantization, but motivated by the unary admissibility constraint rather than by rate-distortion considerations.

In RSVP terms, the TCU encoding defines a non-uniform admissibility metric on the weight manifold: small weights have a fine-grained admissibility structure (many distinguishable states in a small neighborhood) while large weights have a coarser structure. The entropy field S is correspondingly non-uniform, higher for large weights and lower for small weights. Since empirically the most attack-sensitive weights tend to have small magnitudes, the TCU encoding concentrates the low- S region precisely where adversarial vulnerability is highest.

Worked Example

TCU memory overhead for 8-bit VGG-8. For a network where weights are stored in 8-bit BCD, the full unary representation requires $2^8 - 1 = 255$ bits per weight, a $255/8 \approx 32\times$ memory overhead. With TCU encoding, protecting 0.2% of weights (the most sensitive ones) requires only 1.07% additional memory overhead — a $30\times$ reduction. Lu et al. (2024) report that this 0.2% protection rate, combined with post-attack weight locking (Section 21.4), achieves 86.73% mean post-attack accuracy on VGG-8/CIFAR-10, compared to 32.89% without any defense, at a total memory overhead below 3%.

21.4 Post-Attack Recovery via Sensitivity-Aware Weight Locking

Pre-attack unary protection covers only a small fraction α of weights. An adversary with Hamming distance budget HD will eventually attack unprotected weights. Post-attack recovery requires: (1) detection of which weights were flipped, and (2) correction of the flipped values.

Detection uses group-wise MSB checksum verification: weights are partitioned into groups of size G ; a checksum mismatch flags the entire group as potentially compromised. The choice of G trades detection precision (smaller G localizes attacks more precisely) against memory overhead (smaller G requires more checksum storage).

Definition 21.4 (Weight locking). Given K pre-computed *centroids* $\{W_k\}_{k=1}^K$ for a layer’s weights, *weight locking* replaces all weights in a detected victim group with their nearest centroid. Centroids are computed pre-deployment via a sensitivity-aware K -means variant in which the clustering distance is the locking-induced accuracy loss (equation (21.1)) rather than Euclidean distance:

$$d_{in} = \nabla_{W_i} \mathcal{L} \cdot (W_i - \widetilde{W}_n) + \frac{1}{2} \nabla_{W_i}^2 \mathcal{L} \cdot (W_i - \widetilde{W}_n)^2, \quad (21.3)$$

where \widetilde{W}_n is the centroid of the n -th detection group.

Weight locking generalizes the simpler pruning-based recovery method, which forces detected victim weights to zero. Pruning to zero is a special case of weight locking with $K = 1$ and centroid $W_1 = 0$. The locking approach allows $K > 1$ centroids that better approximate the pre-attack weight distribution, recovering higher accuracy at comparable or lower memory cost.

The memory overhead of the locking scheme is

$$m_L = \frac{|W|(\log_2 K + 2)}{G \cdot b|W|} \quad (G > 1), \quad (21.4)$$

where the $\log_2 K$ term stores the cluster ID per weight and the $2/G$ term stores the checksum. The accuracy-constrained optimization over (G, K) is solved offline by a greedy search that initializes at $(G = 512, K = 1)$ and halves G while doubling K until the accuracy drop threshold η is satisfied.

21.5 RSVP Interpretation: Non-Ideality as Admissibility Engineering

The full defense framework of Lu et al. (2024) — TCU pre-attack encoding combined with sensitivity-aware post-attack weight locking — is, in RSVP terms, a deliberate engineering of the admissibility field of the weight manifold \mathcal{M}_W .

The TCU encoding restricts the admissibility radius of the adversary’s action space (Proposition 21.2): it lowers S in the neighborhood of vulnerable weights, making large-perturbation trajectories inaccessible to the attacker. The weight locking pre-assigns centroid attractors in \mathcal{M}_W : after an attack, the KES-like recovery map sends detected victim weights to their pre-designated centroid, analogous to a **Refuse** operation that filters post-attack configurations to those consistent with the pre-deployment constraint set.

Proposition 21.5 (Weight locking as a Refuse operation). *Let $\mathcal{C}_{\text{lock}} = \{W_k\}_{k=1}^K$ be the pre-deployed centroid set. Then weight locking is equivalent to applying $\mathbf{Refuse}(\mathcal{E}_W, \mathcal{C}_{\text{lock}})$: it filters the post-attack weight configuration to the nearest element of $\mathcal{C}_{\text{lock}}$, exactly as \mathbf{Refuse} filters events to those consistent with a constraint set.*

Proof. The \mathbf{Refuse} operator (Definition in Chapter 12) selects events consistent with \mathcal{C} . Here, the events are weight configurations and $\mathcal{C}_{\text{lock}}$ is the constraint that the weight must lie in $\{W_k\}_{k=1}^K$. Locking maps each detected weight to $\arg \min_k \|W_i - W_k\|$, which is the element of $\mathcal{C}_{\text{lock}}$ nearest to the post-attack weight — precisely the nearest consistent configuration. This is a special case of \mathbf{Refuse} under the Euclidean proximity consistency condition. \square

The combined TCU + locking framework therefore implements a two-stage RSVP defense: the first stage (TCU) engineers the entropy field S of \mathcal{M}_W to be low in adversarially vulnerable regions, making large perturbations inadmissible; the second stage (locking) applies a \mathbf{Refuse} -like correction that restores violated configurations to the nearest admissible centroid.

Theorem 21.6 (Admissibility-security duality). *Let \mathcal{M}_W be the weight manifold with RSVP entropy field S_W encoding local admissibility radius. Let HD be an adversary’s Hamming distance budget and let $S_W^{\max}(\theta) = \max_{x \in B_{HD}(\theta)} S_W(x)$ be the maximum entropy accessible to the adversary within Hamming distance HD of weight θ . Then the adversary’s maximum achievable loss perturbation is bounded by*

$$\max_{HD\text{-attack}} |\Delta \mathcal{L}| \leq C \cdot HD \cdot \exp(S_W^{\max}(\theta)), \quad (21.5)$$

for a constant C depending on the second-order structure of \mathcal{L} .

Proof. Each bit-flip of a TCU-encoded weight changes W_i by at most $\Delta_{\text{TCU}} = \exp(-\hat{b}_i)$ in normalized units. The entropy field $S_W(\theta)$ at weight θ is the log-volume of the admissible perturbation neighborhood, which under TCU encoding satisfies $S_W \leq \log(\Delta_{\text{TCU}}) =$

$-\hat{b}_i \log 2$. Since each of the HD flips contributes at most $C' \exp(S_W)$ to the loss perturbation (by the first-order term in (21.1) and $|\Delta W_{\text{TCU}}^{\text{MSB}}| = e^{S_W}$), summing over HD flips gives the bound (21.5). \square

Theorem 21.6 makes the RSVP security principle precise: the adversary’s damage is bounded by the entropy field of the weight manifold times the attack budget. Engineering S_W to be small in sensitive regions is equivalent to engineering adversarial robustness.

21.6 Experimental Grounding

Lu et al. (2024) evaluate the defense framework on VGG-8/CIFAR-10 and ResNet-18/CIFAR-100 with 4-, 6-, and 8-bit quantization under BFA attacks with $HD = 100$. Without any defense, 8-bit VGG-8 achieves 13.52% mean post-attack accuracy (from 88% clean accuracy). The combined TCU + weight locking framework recovers 86.73% mean post-attack accuracy at 2.36% total memory overhead. This represents a practical realization of the admissibility-security duality (Theorem 21.6): by engineering S_W to be low in the 0.2% most sensitive weights via TCU, and applying a **Refuse**-like locking correction post-attack, the framework recovers near-ideal accuracy at marginal storage cost.

The key quantitative relationships that ground the RSVP interpretation are:

Sensitivity reduction TCU reduces MSB-flip perturbation size by $2^{b-1} : 1$ relative to BCD (Proposition 21.2), which reduces the first-order term in sensitivity (21.1) by the same factor.

Memory compression TCU achieves 6–11 \times memory reduction over full unary (Eq. (21.2)), enabling protection of a sufficient fraction α of vulnerable weights within practical overhead.

Post-attack recovery Weight locking with $K > 1$ centroids outperforms pruning ($K = 1$, centroid at zero) at equal or lower

memory overhead, confirming that the **Refuse** operation with a richer constraint set $\mathcal{C}_{\text{lock}}$ provides better admissibility recovery.

Noise as partial defense Hardware Gaussian noise with standard deviation 0.005 provides some protection by adding uncertainty to the adversary’s gradient estimates, but larger noise levels degrade clean accuracy faster than they improve robustness. In RSVP terms, noise increases S_W uniformly (more uncertainty everywhere), which helps the defender but also degrades the prior landscape Φ .

Epistemic Status

The formal identification of TCU encoding with engineering the RSVP entropy field S_W , and of weight locking with a **Refuse** operation, is an interpretation within the RSVP framework rather than a claim made by Lu et al. (2024). Their paper establishes the empirical and algorithmic results; the RSVP formalization is the contribution of this section. Theorem 21.6 is a bound derived within the RSVP framework under assumptions (TCU encoding, first-order loss approximation) that are validated by the experiments of Lu et al. but are stated as approximations in that work.

Chapter 22

Topological Defects in RSVP Fields

Chapter Dependencies

This chapter depends on the RSVP field equations (9.2)–(9.4) and the admissibility framework (Chapter 8).

22.1 Entropy Vortices and Accessibility Horizons

A *topological defect* in the RSVP fields is a point (or manifold) where the field configuration is singular: Φ diverges, $S \rightarrow 0$, or admissibility fibers disconnect. These defects are not numerical artefacts; they correspond to physically meaningful boundaries in the admissibility structure.

The *entropy vortex* is characterized by non-zero circulation $\Gamma = \oint_C \mathbf{v} \cdot d\ell \neq 0$, indicating a persistent loop in the information flow. Non-zero Γ means the admissibility flow does not converge to a fixed point but circulates—a signature of periodic or quasi-periodic

inference dynamics, such as the oscillatory non-Markovian revivals of Theorem 6.2.

An *accessibility horizon* is a region where the RSVP velocity field \mathbf{v} flows outward in all directions: no trajectory can enter from outside. Inside the horizon, the entropy field S may be very low (highly constrained dynamics) or effectively decoupled from the exterior.

Conjecture 22.1. Long-lived physical and cognitive structures correspond to stabilized entropy defects in RSVP field geometry: configurations where $S \approx 0$ in a localized region (strong constraint, persistent memory) surrounded by $S > 0$ (flexible accessibility).

Worked Example

Entropy-vortex simulation. On a 64×64 lattice with periodic boundary conditions, initialize $\Phi = \cos(2\pi x/L) \cos(2\pi y/L)$, $S = 0.5 + 0.3 \sin(2\pi x/L)$, $\mathbf{v} = 0$. Evolve for $T = 200$ steps with $\alpha = 0.1$, $\beta = 0.05$, $\lambda = 0.01$. Entropy vortices form at the saddle points of Φ where \mathbf{v} circulates around the S gradient maximum. Vortex strength (circulation Γ) grows as α increases.

Chapter 23

RSVP Thermodynamics

Chapter Dependencies

This chapter reinterprets thermodynamic quantities using the RSVP field triple, depending on Chapters 8 and the entropy-accessibility identification.

23.1 Temperature as Accessibility Gradient

Define the effective temperature in a RSVP system as

$$T^{-1} = \frac{\partial S}{\partial E}, \quad (23.1)$$

where E is the effective energy of the admissibility configuration. This parallels the thermodynamic definition $T^{-1} = \partial S / \partial E$ but applies to the accessibility entropy field rather than thermodynamic entropy.

23.2 Entropy Production and the Admissibility Dissipation Functional

The entropy production rate in the RSVP framework is

$$\Sigma = \int_{\mathcal{M}} |\nabla S|^2 d\text{vol}_g \geq 0, \quad (23.2)$$

measuring the spatial inhomogeneity of the accessibility field.

Theorem 23.1 (Minimum admissibility dissipation). *Steady RSVP states (solutions to $\partial_t S = 0$) minimize Σ subject to fixed boundary conditions on S .*

Proof. $\partial_t S = 0$ in equation (9.4) gives $\Delta S \propto S$ at leading order (for \mathbf{v} given by (9.3)). This is the Euler-Lagrange equation for the functional $\int |\nabla S|^2$, so steady states minimize the Dirichlet energy $\int |\nabla S|^2$, which equals Σ . \square

Chapter 24

The Geometry of Learning and Semantic Compression

Chapter Dependencies

This chapter connects the RSVP framework to machine learning and representation theory, depending on the admissibility projection (8.1) and the ELBO (3.2).

24.1 Learning as Accessibility Compression

Learning is the process of finding a projection $\pi_\theta : X \rightarrow M$ that preserves the admissibility distinctions relevant to a given task while discarding irrelevant ones. The compression operator is

$$\mathcal{R}_\theta : X \rightarrow X' \tag{24.1}$$

where X' is a reduced trajectory space. Generalization corresponds to the stability of \mathcal{R}_θ under perturbations of the training trajectories: if nearby trajectories in X map to the same region of X' , the model generalizes; if nearby trajectories are separated by \mathcal{R}_θ , the model overfits.

24.2 CLIO: Constraint-Leveraged Inference and Optimization

The CLIO operator acts on the constraint structure rather than on parameters:

$$\text{CLIO} : (X, \mathcal{A}, \pi) \rightarrow (X', \mathcal{A}', \pi'). \quad (24.2)$$

A CLIO update penalizes destructive projection and accessibility collapse:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} [\mathcal{L}(\pi_{\theta}(x), d) + \lambda \mathcal{D}(\pi_{\theta}) + \mu \mathcal{E}(S)], \quad (24.3)$$

where $\mathcal{D}(\pi_{\theta})$ penalizes fiber collapse and $\mathcal{E}(S)$ penalizes entropy field collapse toward zero.

Technical Note

Semantic attractors and transformer embeddings. Stable regions in representational geometry correspond to semantic attractors: configurations of (Φ, S) where the RSVP velocity field \mathbf{v} converges. In transformer models, token embeddings that are stable under attention operations correspond to high- Φ , low- S regions: high salience, low residual uncertainty. The CLIO penalty $\mathcal{E}(S)$ prevents collapse to a single attractor, maintaining diversity of representation.

Chapter 25

Functorial Structure of Spherepop

Chapter Dependencies

This chapter develops the category-theoretic structure of the Spherepop operators, depending on Chapter 12 and requiring familiarity with symmetric monoidal categories.

25.1 Spherepop as a Symmetric Monoidal Category

Define the *Spherepop category* \mathbf{Sp} with: objects: event histories H_t ; morphisms: admissible transitions $H_t \rightarrow H_{t'}$ (causal extensions); tensor product: \mathbf{Bind} (event composition); unit: the null event \perp .

Proposition 25.1 (\mathbf{Sp} is a strict symmetric monoidal category). *The Spherepop category with \mathbf{Bind} as tensor product satisfies the axioms of a strict symmetric monoidal category: associativity of \mathbf{Bind} , existence of unit \perp , and commutativity up to natural isomorphism.*

Proof. Associativity: $\mathbf{Bind}(\mathbf{Bind}(e_1, e_2), e_3) = \mathbf{Bind}(e_1, \mathbf{Bind}(e_2, e_3))$

by the associativity of \sqcup in the event algebra. Unit: $\mathbf{Bind}(e, \perp) = e$ by $e \sqcup \perp = e$. Commutativity up to isomorphism: $\mathbf{Bind}(e_1, e_2) \cong \mathbf{Bind}(e_2, e_1)$ since \sqcup is commutative for causally compatible events. \square

25.2 Collapse as a Functor and Topological Obstructions

Collapse : $\mathcal{E} \rightarrow \mathcal{H}_{\text{eff}}$ is a functor from the event algebra to the category of effective Hilbert spaces, mapping event composition to tensor product of compatible subspaces.

Conjecture 25.2. The non-trivial cohomology classes of the Spherepop category \mathbf{Sp} correspond to irreversible physical memory structures: configurations where the holonomy of the admissibility connection (Chapter 18) is non-trivial.

Chapter 26

Computational Complexity of RSVP Inference

Chapter Dependencies

This chapter estimates scaling behavior for RSVP field evolution and KES dynamics, depending on the field equations (9.2)–(9.4) and the KES memory kernel (13.4).

26.1 Scaling of Field Evolution

Theorem 26.1 (Single-step RSVP complexity). *Under bounded entropy gradients $|\nabla S| \leq C$, a single step of RSVP field evolution on an N -cell lattice (using finite differences) is $O(N)$.*

Proof. Each lattice cell requires $O(1)$ operations: computing the Laplacian of Φ (via $O(d)$ neighbor values in d dimensions), updating \mathbf{v} from (9.3), and updating S from (9.4). The total is $O(dN) = O(N)$ for fixed d . The entropy gradient bound $|\nabla S| \leq C$ ensures the nonlinear $\log S$ term does not cause blow-up. \square

26.2 Memory Cost and Approximate Inference

The KES memory kernel K_m of depth m requires storing the last m event vectors: $O(m|\mathcal{E}|)$ memory. For deep memory ($m \gg 1$), variational RSVP approximates the full kernel by a parametric family:

$$K_m^\phi(e_{t-m}, \dots, e_{t-1}) \approx f_\phi(e_{t-m}, \dots, e_{t-1}), \quad (26.1)$$

where f_ϕ is a neural network trained to minimize the KL divergence from the true kernel. This gives $O(|\phi|)$ memory independent of m , at the cost of approximation error.

Technical Note

GPU parallelism for RSVP evolution. Each RSVP lattice cell is an independent CUDA thread. Fields $(\Phi_i, \mathbf{v}_i, S_i)$ are stored as structure-of-arrays in global memory; the update kernel (9.2)–(9.4) is launched with N threads per step. Asynchronous entropy propagation uses shared memory for the $\text{div}(S\mathbf{v})$ term within each warp.

Complexity summary.

Component	Per-step cost	Memory
RSVP field evolution	$O(N)$	$O(N)$
KES event selection	$O(\mathcal{E})$	$O(\mathcal{X}_t)$
KES memory (m steps)	$O(m \mathcal{E})$	$O(m \mathcal{E})$
Variational KES	$O(\phi)$	$O(\phi)$
TARTAN (K cells)	$O(KN)$	$O(KN)$

Chapter 27

RSVP Cosmological Optics

Chapter Dependencies

This chapter applies the RSVP propagator (11.1) to cosmological photon propagation. It is conjectural throughout; epistemic status is stated explicitly.

Epistemic Status

This chapter presents conjectural cosmological applications of RSVP rather than experimentally established physics. All results should be understood as predictions of the framework contingent on its physical identification with the photon propagation equation.

27.1 Entropic Redshift

The RSVP propagator predicts a cosmological photon frequency shift

$$\frac{\Delta\omega}{\omega_0} = \frac{1}{2}(\Phi(x_{\text{obs}}) - \Phi(x_{\text{emit}})) - \frac{\alpha}{2} \int_{\text{path}} S ds. \quad (27.1)$$

If Φ decreases along the line of sight (earlier epochs have lower likelihood potential), this produces a redshift without metric expansion. The entropic attenuation term modifies the luminosity-distance relation relative to Λ CDM.

27.2 Structure Formation as Accessibility Condensation

Large-scale structure forms where the RSVP scalar field Φ develops local maxima (high admissibility basins): matter and energy are drawn toward high- Φ regions by the RSVP velocity field $\mathbf{v} = (\alpha/S)\nabla\Phi$. This replaces gravitational collapse as the primary mechanism of structure formation, with the gravitational potential identified as a component of Φ .

Conjecture 27.1. Observed cosmological acceleration may arise from large-scale entropy-gradient effects ($\nabla S \neq 0$ on Hubble scales) rather than from a cosmological constant.

Part VII

Toward a Unified Accessibility Calculus

Chapter 28

Projection, Agency, and Observer Geometry

Chapter Dependencies

This chapter depends on the admissibility framework (Chapter 8) and the Sphero-pop axioms (Chapter 12).

28.1 Observers as Recursive Projection Systems

An observer is a system that recursively projects X onto a sequence of increasingly compressed representations $M_1 \supset M_2 \supset \dots$, retaining only admissible distinctions at each level. Define the *observer projection map* $\pi_i : X \rightarrow M_i$, where M_i is the i -th observer's representational manifold.

Definition 28.1 (Consensus manifold). The *consensus manifold* for a set of observers $\{\pi_i\}$ is

$$M_{\text{consensus}} = \bigcap_i \pi_i(X), \quad (28.1)$$

the intersection of representational images restricted to mutually stable admissibility regions.

Shared classical reality corresponds to a large, stable consensus manifold. Quantum non-classicality corresponds to configurations where different observer projections disagree: $\pi_i(x) \neq \pi_j(x)$ for trajectories x in the non-classical regime.

Conjecture 28.2. Persistent observer identity corresponds to a stable fixed point of recursive admissibility compression: a configuration π^* such that $\pi^* \circ \pi^* \approx \pi^*$ in the sense that repeated self-projection does not substantially reduce the admissibility structure.

Chapter 29

The Unified Accessibility Calculus and Open Problems

Chapter Dependencies

This final chapter synthesizes all preceding material and defines the long-term research program. All results marked as conjectures or open problems are formally stated but unproved.

29.1 The Universal Accessibility Functional

Define the *total admissibility* of a system at time t :

$$\mathfrak{A}(t) = \int_X \mathcal{A}(x, t) dx. \quad (29.1)$$

The unified evolution equation for the admissibility functional is proposed as

$$\partial_t \mathcal{A} = \mathcal{G}(\Phi, \mathbf{v}, S, H_t), \quad (29.2)$$

where \mathcal{G} is a generalized admissibility operator encoding the coupled RSVP dynamics, the KES event selection, and the history-dependent

memory kernel. The RSVP field equations (9.2)–(9.4) are the leading-order approximation to \mathcal{G} in the limit of small history dependence ($m = 0$) and compact manifold.

29.2 Mathematical Open Problems

Global existence. Theorem 9.2 gives local-in-time well-posedness. Whether solutions exist globally or develop finite-time singularities (entropy field reaching zero) remains open. The lower bound $S \geq s_0 e^{-Ct}$ suggests sufficiently strong β is needed.

Categorical coherence of \mathbf{Sp} . The Spherepop category \mathbf{Sp} should admit a dagger structure encoding time-reversal asymmetry. Whether \mathbf{Sp} embeds into the Abramsky-Coecke categorical framework for quantum mechanics is open.

Fiber bundle curvature and Berry phase. The conjecture that RSVP admissibility connection holonomy equals the quantum mechanical Berry phase requires a formal identification of the RSVP horizontal distribution with the quantum geometric tensor.

Quantization of RSVP. The classical field theory $\mathcal{L}_{\text{RSVP}}$ should be quantized. The positivity of S in the classical theory must be preserved as a constraint on the quantum Hilbert space.

Complexity of exact KES. The exact KES map requires solving the optimization $\max_e \mathcal{F}(e)$ over \mathcal{E} . The computational complexity of this optimization as a function of $|\mathcal{E}|$ and $|\mathcal{X}_t|$ is not known in general.

29.3 The Four Core Ideas

Every open problem, extension, and application in this text is anchored to the same four conceptual pillars, which the reader should

hold in view throughout:

1. **Admissibility:** not all configurations are equally accessible; differential accessibility is the true primitive.
2. **Projection:** the map $\pi : X \rightarrow M$ and its fiber structure encode the information cost of conventional modeling.
3. **Entropy as accessibility:** S measures future trajectory volume, not merely Shannon uncertainty.
4. **Irreversible event structure:** events are primitive; states are equivalence classes of event histories; irreversibility is foundational, not emergent.

These four ideas are the invariants of the framework. Any section that cannot be connected to at least two of them is a candidate for removal.

*CHAPTER 29. THE UNIFIED ACCESSIBILITY CALCULUS AND
OPEN PROBLEMS*

Bibliography

- Amari, S. (2016). *Information Geometry and Its Applications*. Springer.
- Lu, H., Yin, Z., Bhoumik, P., Banerjee, S., Chakrabarty, K., and Gu, J. (2024). The unlikely hero: Nonideality in analog photonic neural networks as built-in defender against adversarial attacks. *arXiv:2410.01289*. To appear in *Proc. ASP-DAC*, 2025.
- Born, M. (1926). Zur Quantenmechanik der Stossvorgänge. *Zeitschrift für Physik*, 37, 863–867.
- Caticha, A. (2019). The entropic dynamics approach to quantum mechanics. *Entropy*, 21(10), 943.
- Fuchs, C. A., Mermin, N. D., and Schack, R. (2014). An introduction to QBism. *American Journal of Physics*, 82(8), 749–754.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. 3rd ed. CRC Press.
- Gleason, A. M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6(6), 885–893.

- Gorini, V., Kossakowski, A., and Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N -level systems. *Journal of Mathematical Physics*, 17(5), 821–825.
- Griffiths, R. B. (1984). Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics*, 36(1), 219–272.
- Hong, C. K., Ou, Z. Y., and Mandel, L. (1987). Measurement of subpicosecond time intervals between two photons by interference. *Physical Review Letters*, 59(18), 2044.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106(4), 620.
- Kasap, S. O. (2013). *Optoelectronics and Photonics: Principles and Practices*. 2nd ed. Pearson.
- Lindblad, G. (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2), 119–130.
- Mandel, L. and Wolf, E. (1995). *Optical Coherence and Quantum Optics*. Cambridge University Press.
- Rivas, A. and Huelga, S. F. (2012). *Open Quantum Systems: An Introduction*. Springer.
- Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics*, 35(8), 1637–1678.
- Schlosshauer, M. (2007). *Decoherence and the Quantum-to-Classical Transition*. Springer.
- Soshnikov, A. (1999). Universality at the edge of the spectrum in Wigner random matrices. *Communications in Mathematical Physics*, 207(3), 697–733.
- Tracy, C. A. and Widom, H. (1994). Level-spacing distributions and the Airy kernel. *Communications in Mathematical Physics*, 159(1), 151–174.

Wakefield, J. (2013). *Bayesian and Frequentist Regression Methods*. Springer.

Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3), 715.