Three Admissibility Classes and Maxwell-Form Persistent Transport

A Synthesis Bridging the σ-Sector Master-Action Framework, the Hamiltonian Admissibility Derivation, and the Maxwell-Form Uniqueness Theorem

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The VERSF programme has, in three separate strands of work, established three results that have so far been pursued in parallel rather than as one architecture.

The first is the σ-sector master-action framework, which describes how the substrate restores admissibility when it has been pushed off the admissible manifold. This is a *dissipative* dynamics — it smooths things out, dissipates structure, and drives the substrate back toward equilibrium. The σ-family papers derived this dynamics in detail, with the σ-sector emerging as the natural admissibility-restoring response of the $K = 7$ architecture.

The second is the Hamiltonian admissibility derivation, which shows that any evolution that is composable, continuous, and distinguishability-preserving must take the form of unitary evolution generated by a self-adjoint Hamiltonian operator. This is the *reversible* dynamics that quantum mechanics uses — and the derivation establishes that the Hamiltonian is not an independent postulate but a structural consequence of the reversibility requirements themselves, derived via Wigner's theorem and Stone's theorem.

The third is the Maxwell admissibility uniqueness theorem, which establishes that under the substrate axioms BCB (Bit Conservation and Balance) and TPB (Ticks-Per-Bit), within a specific admissibility class motivated by locality, gauge redundancy, and closure-geometry covariance, the unique surviving continuum dynamics is Maxwell-form U(1) gauge transport. This is the structure of classical electromagnetism — not postulated but derived as the unique admissible dynamics on the persistent observable sector.

This paper integrates these three strands into a unified architecture organised around the recognition that they each describe a structurally distinct *admissibility class*, with its own characteristic dynamics, its own appropriate variational structure, and its own role in the broader VERSF programme.

The σ-sector belongs to the *dissipative admissibility class* — admissibility-restoring dynamics governed by gradient flow. The persistent gauge sector (the topologically protected directions surviving under refinement) belongs to the *reversible admissibility class* — distinguishability-preserving transport governed by unitary Hamiltonian evolution. The record-formation sector belongs to a third *irreversible admissibility class* — commitment dynamics governed by GKSL-type Lindbladian evolution. Each class is a structurally distinct kind of admissibility, each has its own natural variational structure, and each contributes a different aspect of the physical content of the broader theory.

The substantive new content of this paper is the integration of these three classes into one architecture, and the establishment of a synthesis theorem that composes the three strands to identify the natural continuum-limit dynamics on the persistent sector. The synthesis runs:

- The σ-sector master-action framework establishes the persistent gauge sector as $\mathrm{Ker}(\delta^*) \cong H^1$ — the topologically protected refinement-stable observable sector.
- The Hamiltonian admissibility derivation establishes that reversible distinguishability-preserving evolution on any admissible system necessarily takes unitary Hamiltonian form $U(\tilde{\tau}) = e^{(-iH\tilde{\tau})}$.
- The Maxwell admissibility theorem establishes that under BCB/TPB and the admissibility class (B1)–(B4), the unique continuum-limit dynamics on H^1 is Maxwell-form U(1) gauge transport.

Composing these three: under BCB/TPB and the admissibility class (B1)–(B4), the σ-sector framework's persistent gauge sector — considered in its continuum limit as a refinement-stable observable sector — carries Maxwell-form U(1) gauge transport as the unique admissible reversible dynamics. This is the central synthesis theorem of the paper.

The paper does *not* claim to derive the full machinery of conventional electrodynamics — full relativistic completion, matter coupling, quantisation, coupling constants, and phenomenological embedding remain open. What it claims is narrower but conceptually substantial: that under the substrate axioms of the broader VERSF programme, the three admissibility classes form a coherent architecture, and within that architecture the persistent gauge sector of the σ-sector framework carries Maxwell-form U(1) transport as the unique refinement-stable reversible dynamics on the persistent observable sector.

The paper is therefore a *bridge paper*: it does not derive its component strands (they are derived in their respective companion papers) but it shows that those strands compose into a unified architecture, and that within that architecture a substantively positive synthesis theorem holds. The σ-sector / persistent gauge sector unification is, in this synthesis, given a positive resolution: the σ-sector and the persistent gauge sector belong to structurally distinct admissibility classes, and the master-action unification is the recognition that they couple consistently within the three-class architecture, with the persistent sector carrying Maxwell-form U(1) transport as its natural reversible dynamics.

Abstract

The σ-sector master-action framework paper established the kinematic structure of the master configuration space $C^1(C_6)$ on the rim cycle of the K = 7 architecture, with the orthogonal Hodge decomposition into the σ-sector subspace Im(δ) and the persistent gauge sector subspace Ker(δ*) ≅ $H^1(C_6)$. Under uniform gradient-flow treatment, the σ-sector exhibited its expected dissipative dynamics but the persistent gauge sector was left static. The framework paper identified the upgrade to non-trivial dynamics on Ker(δ*) as the load-bearing open problem.

This paper resolves the open problem by recognising that the σ-sector and the persistent gauge sector belong to *structurally distinct admissibility classes*, each with its own appropriate variational structure, and by composing this recognition with two established results elsewhere in the VERSF programme — the Hamiltonian admissibility derivation and the Maxwell admissibility uniqueness theorem — to identify Maxwell-form U(1) gauge transport as the unique admissible continuum dynamics on the persistent sector.

We establish four theorems, organised into a four-part architecture:

Part I — Dissipative admissibility (σ-sector). *Lemma 1 (Classification of gradient-flow dynamics).* Gradient-flow variational principles are intrinsically first-order in time, generating either invariant or strictly dissipative dynamics. Oscillation is structurally excluded. This classical fact, applied to the dissipative admissibility class, characterises the σ-sector's structural shape.

Theorem 2 (K = 7 catalogue does not generate mass-like terms). The constraint catalogue extended from spoke variables to 1-cochains admits the same four-element structure with no mass-like terms. This is a feature of the dissipative class, not a deficiency: it does not preclude Hamiltonian structure on the persistent sector, which belongs to a different admissibility class.

Part II — Reversible admissibility (persistent gauge sector). *Theorem 3 (Reversible admissibility on the persistent sector).* The persistent gauge sector Ker(δ*) is the natural carrier of reversible admissibility within the master configuration space. Under the Hamiltonian admissibility derivation (Wigner + Stone), reversible distinguishability-preserving evolution on Ker(δ*) necessarily takes unitary Hamiltonian form $U(\tilde{\tau}) = e^{(-iH_h \tilde{\tau})}$ with self-adjoint H_h, preserving the Wilson-loop topological invariant exactly.

Part III — Persistent cohomological transport (Maxwell-form). *Main Synthesis Theorem (Maxwell-form persistent transport).* The persistent gauge sector Ker(δ*) ≅ $H^1(C_6)$, considered as the K = 7 specialisation of the general substrate's refinement-persistent cohomological sector $H^1(G(\Lambda))$, satisfies the conditions of the Strong Synthesis Theorem (Maxwell admissibility + persistent sector identification). Under BCB, TPB, and the admissibility class (B1)–(B4), the unique refinement-stable continuum-limit reversible dynamics on the persistent sector is Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$.

Part IV — Unified three-class admissibility architecture. *Theorem 4 (Three admissibility classes generate three variational structures).* The dissipative class (σ-sector, gradient flow), the

reversible class (persistent sector, unitary Hamiltonian), and the irreversible record class (commitment formation, GKSL Lindbladian) form a unified three-class admissibility architecture. The σ-sector / persistent gauge sector unification proceeds via inter-class coupling within this architecture, with matter coupling providing observable propagation and effective inertia.

The Main Synthesis Theorem is the central new result of this paper. It is established not by new derivation but by composition of three results derived in their respective companion papers: the σ-sector master-action framework, the Hamiltonian admissibility derivation, and the Maxwell admissibility theorem (with its Strong Synthesis extension). The composition is non-trivial — it requires the sector-classification reframe that recognises the persistent gauge sector as belonging to the reversible admissibility class — and it gives a positive resolution to the framework paper's load-bearing open problem.

Epistemic status. *Proven (within this paper):* Lemma 1 and Theorems 2–4 above; the sector-classification of σ-sector as dissipative admissibility, persistent gauge sector as reversible admissibility, record-formation sector as irreversible admissibility; the three-class architecture and its consistency; the Main Synthesis Theorem as a composition of established results. *Composed from companion papers (load-bearingly cited):* the Hamiltonian admissibility derivation (Wigner + Stone derivation of $U(t) = e^{-iHt}$ from axioms A1 composability, A2 continuity, A3 distinguishability preservation); the Maxwell admissibility uniqueness theorem (BCB/TPB + (B1)–(B4) → Maxwell-form U(1)); the Strong Synthesis Theorem (Maxwell-form on $H^1(G(\Lambda))$ as unique refinement-stable continuum dynamics). *Postulated (structural inputs to the sector-classification):* the assignment of σ-sector to the dissipative admissibility class, persistent gauge sector to the reversible admissibility class, and record-formation sector to the irreversible admissibility class. *Open:* the specific eigenvalue ω of the persistent-sector generator ($\omega = 0$ valid, $\omega > 0$ valid; pinning down ω requires inter-class coupling or matter coupling); full relativistic completion of the persistent sector dynamics beyond Maxwell-form at $O(\varepsilon^0)$; matter coupling derivation and observable propagation; quantisation structure; coupling constants and phenomenological embedding; higher-dimensional lifting of the K = 7 architecture; explicit derivation of how the K = 7 substrate embeds in the more general substrate of the Maxwell admissibility theorem.

Contents

PART I — DISSIPATIVE ADMISSIBILITY

1. Introduction

The σ-sector master-action framework paper established the cohomological structure underlying the σ-sector / persistent gauge sector unification on the rim cycle of the K = 7 architecture: the master configuration space $C^1(C_6)$, the orthogonal Hodge decomposition into Im(δ) (σ-sector) and Ker(δ*) (persistent gauge sector), and the identification of the persistent direction simultaneously as the kernel of the Hodge Laplacian and as the first cohomology class $H^1(C_6) \cong \mathbb{R}$.

The framework paper treated both sectors using a single variational structure: the gradient-flow form inherited from the K = 7 master-action variation. Under this uniform treatment, the σ-sector exhibited its expected dissipative dynamics, but the persistent gauge sector was left static. The framework paper identified the upgrade of the persistent sector to non-trivial dynamics as the load-bearing open problem (P1).

This paper resolves the open problem in two stages, each requiring a different kind of structural input.

The first stage is the *sector-classification* recognition that the σ-sector and the persistent gauge sector are most naturally assigned to two structurally distinct admissibility classes — each with its own appropriate variational structure. The K = 7 master-action gradient-flow principle is the natural variational structure for the *dissipative admissibility class*, of which the σ-sector is the natural carrier. The persistent gauge sector is naturally assigned to a different admissibility class, the *reversible admissibility class*, with a different appropriate variational structure (unitary Hamiltonian evolution).

The framework paper's static persistent sector was *not* a misapplication or error: the framework paper explicitly flagged (§5.1) that the Hodge heat flow everywhere on $C^1(C_6)$ was one specific postulate among several, with the alternative — "postulating a Hamiltonian operator on Ker(δ*) alongside the dissipative Hodge flow on Im(δ)" — explicitly noted as available. The framework paper had the architectural shape right; what it didn't have was the *structural source* of the Hamiltonian alternative. The present paper makes the alternative substantive by sourcing it from the reversible admissibility class — drawing on the Hamiltonian admissibility derivation work elsewhere in the broader programme. The contribution is the sharpening of the framework paper's explicitly-acknowledged postulate via a structural source, not a correction of an error. This is the substantive content of Parts I and II of the present paper.

The second stage is the *composition with established results* in the broader VERSF programme. Two strands of established work compose with the sector-classification to give a substantively positive resolution. The first is the Hamiltonian admissibility derivation, which establishes via Wigner's theorem and Stone's theorem that any admissible reversible distinguishability-preserving evolution necessarily takes unitary form $U(t) = e^{-iHt}$ with a self-adjoint generator H. Applied to the persistent gauge sector, this establishes Hamiltonian structure as the necessary form of reversible-class dynamics on Ker(δ*). The second is the Maxwell admissibility uniqueness theorem, which establishes that under the substrate axioms BCB and TPB, within an admissibility class (B1)–(B4) motivated by locality, gauge redundancy, and closure-geometry covariance, the unique continuum-limit dynamics on the persistent observable sector is Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$. Combined with the Strong Synthesis Theorem (which identifies $H^1(G(\Lambda))$ as the unique refinement-stable observable sector), this gives the positive synthesis: under BCB/TPB and (B1)–(B4), the persistent gauge sector of the σ-sector framework — considered as the K = 7 specialisation of $H^1(G(\Lambda))$ — carries Maxwell-form U(1) gauge transport as the unique admissible reversible dynamics. This is the substantive content of Part III.

Part IV completes the architecture by integrating a third admissibility class — the *irreversible record-formation class* — corresponding to the commitment / decoherence dynamics of the broader VERSF programme. This class is dynamically distinct from both gradient flow (dissipative class) and unitary evolution (reversible class) — it involves non-unitary state collapse via record formation — and is governed by GKSL-type Lindbladian dynamics. The three classes together — dissipative, reversible, irreversible — form a unified admissibility architecture in which each major sector of the VERSF programme has its place.

The paper is organised by these four parts plus closing sections.

Part I (§§1–7) establishes the dissipative admissibility class and its variational structure. §1 is this introduction. §2 recapitulates the framework paper's uniform variational treatment and identifies the implicit assumption that one variational structure governs both sectors. §3 introduces the dissipative admissibility class. §4 develops gradient flow as its natural variational structure. §5 enumerates admissible quadratic functionals on 1-cochains and establishes the K = 7 catalogue structure. §6 establishes Lemma 1 (classification of gradient-flow dynamics) and §7 proves Theorem 2 (the K = 7 catalogue does not generate mass-like terms).

Part II (§§8–10) establishes the reversible admissibility class and its variational structure. §8 introduces the class. §9 summarises the Hamiltonian admissibility derivation (Wigner + Stone) from the companion paper. §10 proves Theorem 3, establishing reversible admissibility on the persistent gauge sector with Hamiltonian generator structure.

Part III (§§11–15) composes the sector-classification with the substrate-level Maxwell admissibility theorem. §11 develops the Wilson/cohomology identification linking the persistent gauge sector of the σ-sector framework to the broader substrate's refinement-stable observable sector. §12 summarises the Maxwell admissibility uniqueness theorem from the companion paper. §13 summarises the Strong Synthesis Theorem integrating the substrate-level cohomology with the Maxwell admissibility result. §14 states and proves the Main Synthesis Theorem of this paper, composing the three strands to identify Maxwell-form U(1) transport as the unique admissible continuum dynamics on the persistent sector. §15 integrates two further companion papers (the Wilson Limit paper and the Closure-Symmetry paper) to establish substrate-level mechanisms supporting Lorentz-compatible Maxwell-form transport, strengthening the Main Synthesis Theorem from "Maxwell-form structure under Lorentz-covariance admissibility postulate" to "Maxwell-form structure with substrate-level mechanisms for Lorentz-compatibility."

Part IV (§§16–18) completes the unified architecture. §16 proves Theorem 4, integrating the three admissibility classes. §17 introduces the third (irreversible) class. §18 presents the unified architecture as a whole.

Closing sections (§§19–23) discuss the refined inertia distinction, implications for the unification programme, scope, open problems, and conclusion.

---

2. The Framework Paper's Uniform Variational Treatment

The framework paper §5.1 distinguished two postulates underlying the master dynamics:

- **Postulate 1.** The master configuration space is $C^1(C_6)$, the 1-cochains on the rim cycle. This extends the σ-sector envelope configuration space $C^0(C_6)$ to include the persistent gauge sector subspace $\mathrm{Ker}(\delta^*) \subset C^1(C_6)$.
- **Postulate 2.** The master dynamics is the Hodge heat flow $\partial_{\tilde{\tau}} \alpha = -\Delta^1 \alpha$ everywhere on $C^1(C_6)$ — including on $\mathrm{Ker}(\delta^*)$, where Δ^1 acts as zero, giving trivial invariant dynamics on the persistent sector.

The framework paper acknowledged Postulate 2 as a specific choice within a family of possible dynamics. Implicit in both postulates, however, was a deeper structural assumption: that *one variational principle* (gradient flow) governs both the σ-sector and the persistent gauge sector uniformly across the master configuration space. The framework paper's static persistent sector followed directly from this uniform treatment: gradient flow applied to a subspace with no quadratic-potential gradient (the harmonic subspace) gives no dynamics.

The load-bearing implicit assumption was therefore not Postulate 2 itself but the uniformity assumption underlying it. Different postulates within the gradient-flow family give either invariance or dissipation, but never the conservative oscillation that conventional persistent gauge sectors exhibit. The present paper proposes an alternative to the uniformity assumption: that the σ-sector and the persistent gauge sector are most naturally assigned to structurally distinct admissibility classes, each with its own appropriate variational structure. This is not a correction of an error in the framework paper — which explicitly flagged the uniform treatment as a postulate, with the Hamiltonian alternative on $\mathrm{Ker}(\delta^*)$ noted as available — but a sharpening of that postulate via a structural source in the broader VERSF programme.

Postulate 1 (the master configuration space as $C^1(C_6)$) is retained throughout — the cochain-complex setting is the correct kinematic structure, and the present paper builds on it. What this paper changes is the uniform variational treatment underlying Postulate 2: instead of one variational principle governing both sectors, we recognise distinct admissibility classes, each with its own appropriate variational structure.

---

3. The Dissipative Admissibility Class

Dissipative admissibility describes dynamics that move the substrate back toward admissibility when it has been pushed off the admissible manifold. The substrate has strayed (by perturbation, by sub-threshold structure, by constraint violation); admissibility restoration smooths out the deviation, dissipating off-admissible structure, and bringing the configuration back toward the admissibility-satisfying region of state space.

The characteristic features of dissipative admissibility:

- **First-order in time.** The time evolution is determined by the configuration's current location relative to admissibility, not by any history-dependent or higher-order temporal structure.
- **Gradient-flow variational principle.** The dynamics is $\partial_\tau \phi = -\, \partial S[\phi] / \partial\phi$ for an admissibility-violation action S that is minimised on the admissible manifold.
- **Entropy-reducing / structure-dissipating.** Off-admissibility content decreases monotonically along the flow.
- **Convergent.** The flow asymptotically converges to the admissible manifold; once admissible, the dynamics is invariant.

Epistemic flag. The assignment of each VERSF sector to a particular admissibility class is itself a structural postulate, not derived within this paper. We are *postulating* — well-motivatedly but as a postulate — that the σ-sector belongs to the dissipative admissibility class, that the persistent gauge sector belongs to the reversible admissibility class, and that the record-formation sector belongs to the irreversible admissibility class. Each assignment has supporting motivation (the σ-sector's derivation from the K = 7 master-action variation gives it gradient-flow structure by construction; the persistent gauge sector's topological-protection structure means no admissible substrate constraint can dissipate it, which makes the reversible-class assignment natural; the record-formation sector's non-unitary irreversible character is established elsewhere in the programme), but the assignments themselves are structural inputs to the present paper rather than its derived content.

This matters particularly for the persistent gauge sector. The framework paper's §5.1 epistemic note already flagged the dynamics on Ker(δ*) as a postulate: "a different choice — postulating a Hamiltonian operator on Ker(δ*) alongside the dissipative Hodge flow on Im(δ) — would give the same σ-sector dissipation but non-trivial gauge dynamics." The present paper's contribution is not to *derive* that this alternative choice is the right one, but to *source* it from the reversible admissibility class — making the alternative substantive by linking it to the Hamiltonian admissibility derivation work (§9). The structural source is what's new; the assignment itself is a sharpened version of an explicit framework-paper postulate.

The remainder of the paper develops this two-class (and ultimately three-class) framing rigorously, with the assignments of specific VERSF sectors to specific admissibility classes treated as structural postulates throughout.

---

4. The Gradient-Flow Variational Principle

The σ-sector papers derived the on-sector dynamics $\partial_\tau \lambda = -\nabla A_{cl}(\lambda)$ on the spoke variable $\lambda \in \mathbb{R}^6$, where $A_{cl} = \alpha A_{circ} + \gamma A_{comp}$ is the leading-order admissible quadratic action under the constrained-EFT principles (P1a)+(P1b), (P2)–(P4), (P5′). The dynamics is a *gradient flow*: the time derivative of the configuration equals minus the functional gradient of the action.

This gradient-flow structure is *a* natural variational principle of the dissipative admissibility class — there are others (Onsager reciprocal relations, stochastic Langevin dynamics with damping plus noise, etc.) — and is the *specific* variational principle derived within the K = 7 master-action framework. We use it throughout because it is the structure of the σ-sector papers' construction; the narrower correct statement is that gradient flow on the K = 7 catalogue is the dissipative-class variational structure inherited from the σ-sector framework, not that gradient flow is the unique variational structure of dissipative admissibility in general. Specifically:

$$\partial_\tau \phi = -\, \partial S[\phi] \,/\, \partial\phi$$

generates dynamics that decreases S monotonically (since $dS/dt = -\,\|\partial S/\partial\phi\|^2 \leq 0$). For S an admissibility-violation action minimised on the admissible manifold, the gradient flow drives ϕ toward the admissible manifold, dissipating off-admissibility content. The action S functions as a Lyapunov functional for the flow: monotonically decreasing, bounded below, vanishing on the admissible target.

This is exactly the admissibility-restoration dynamics that the σ-sector implements: the substrate is driven by the gradient of the closure-response functional A_{cl} toward configurations where A_{cl} is minimised (the admissible spoke configurations).

The K = 7 master-action variational principle is therefore correctly understood not as a universal variational principle that should also generate the reversible admissibility class's dynamics, but as the natural variational principle of the dissipative admissibility class. Extending it uniformly to the persistent gauge sector — which is naturally assigned to a different admissibility class — gives the static persistent sector of the framework paper; the present paper proposes the alternative assignment (persistent gauge sector → reversible admissibility class), under which the persistent sector carries its own appropriate variational structure (Theorem 3, §10).

Why gradient flow is first-order in time. The gradient-flow equation $\partial_\tau \phi = -\,\partial S/\partial\phi$ is intrinsically *first-order* in time: the configuration's time derivative is determined by the configuration itself (via $\partial S/\partial\phi$), with no second time derivative appearing. This first-order character is a structural feature of the dissipative admissibility class: the class's defining feature is admissibility *restoration* — directional motion toward the admissible manifold — which is naturally first-order, with the current rate of motion determined by how far the configuration currently is from admissibility.

It also has a key consequence: the dynamics generated by gradient flow is either *invariant* (if $\partial S/\partial\phi = 0$ at ϕ, in which case ϕ is a critical point of S and the dynamics fixes it) or *dissipative* (if $\partial S/\partial\phi \neq 0$, in which case ϕ decreases S monotonically toward a critical point). There is no third possibility within gradient flow. *Oscillatory* dynamics requires a second-order time derivative, which the gradient-flow equation does not contain. To get oscillation, one needs a different variational principle — typically a Lagrangian principle with second-order Euler–Lagrange equations — which is, structurally, the variational principle of the reversible admissibility class (§8). The reversible class's Hamiltonian variational structure has built-in second-order temporal content; the dissipative class's gradient-flow structure does not.

---

5. Admissible Quadratic Functionals on 1-Cochains and the K = 7 Catalogue

5.1 The 1-cochain configuration space and its symmetry

The 1-cochain configuration space on the rim cycle C_6 is

$C^1(C_6) = \{ \alpha : E(C_6) \to \mathbb{R} \} \cong \mathbb{R}^6,$

with $E(C_6) = \{e_0, \ldots, e_5\}$ the six rim edges. The dihedral group D_6 of order 12 acts on $C^1(C_6)$ by permuting the edge indices: rotations cyclically permute $e_i \to e_{i+1}$; reflections reverse the cyclic order.

The constrained-EFT principles from the σ-sector papers, extended to 1-cochains:

- **(P1a)** Quadratic at leading order in α.
- **(P1b)** Local — the action depends on values of α at nearby edges, not on arbitrary global integrals (with the exception of topological invariants permitted by the constraint catalogue).
- **(P2)** D_6-covariant.
- **(P3)** Smoothly dependent on substrate structural inputs.
- **(P4)** Integrality.
- **(P5′)** Constraint catalogue closure: the admissible action is a sum of constraint-violation penalties from a closed catalogue of substrate constraints.

5.2 Direct enumeration of candidate quadratic functionals

The natural D_6-invariant quadratic functionals on $C^1(C_6)$ are linear combinations of:

- $\mathbf{A_{mag}^{(1)}(\alpha) := \Sigma_i\, \alpha(e_i)^2}$ — *Mass term.* Does not vanish on Ker(δ*).
- $\mathbf{A_{comp}^{(1)}(\alpha) := \Sigma_i\, (\alpha(e_i) - \alpha(e_{i+1}))^2 = \langle \alpha, \Delta^1 \alpha \rangle}$ — *Edge-difference (Hodge Laplacian).* Vanishes on Ker(δ*).
- $\mathbf{A_{anti}^{(1)}(\alpha) := \Sigma_i\, (\alpha(e_i) + \alpha(e_{i+1}))^2}$ — *Edge-sum.* Does not vanish on Ker(δ*).
- $\mathbf{A_{circ}^{(1)}(\alpha) := (\Sigma_i\, \alpha(e_i))^2 = (\oint_C \alpha)^2}$ — *Circulation conservation.* Does not vanish on Ker(δ*).
- $\mathbf{A_{d2}^{(1)}(\alpha) := \Sigma_i\, (\alpha(e_{i-1}) - 2\alpha(e_i) + \alpha(e_{i+1}))^2 = \langle \alpha, (\Delta^1)^2 \alpha \rangle}$ — *Second-derivative.* Vanishes on Ker(δ*).

5.3 The 1-cochain constraint catalogue

For 1-cochains, the constraint catalogue inherits the four-element structure from the $K = 7$ architecture's substrate constraints, translated from spoke variables to edge 1-cochains. Each catalogue *constraint* corresponds to a substrate admissibility condition; what enters the action A_cl is the quadratic *penalty term* for violation of that constraint. We list the constraint together with its penalty term:

- **A_inc^(1)** — Closure-incidence constraint on 1-cochains. Penalty term off-sector; vanishes on admissible 1-cochains.
- **A_hub^(1)** — Hub anchoring constraint on 1-cochains. Penalty term off-sector; trivially satisfied (no hub variable in 1-cochains on the rim).
- **A_circ^(1)** — Circulation conservation constraint (the constraint is $\oint_C \alpha = 0$). Penalty term: $A_{circ}^{(1)}(\alpha) = (\oint_C \alpha)^2$, which is the squared violation of the constraint.
- **A_comp^(1)** — Edge-competition constraint (penalising differences between adjacent edges). Penalty term: $A_{comp}^{(1)}(\alpha) = \langle \alpha, \Delta^1\alpha \rangle$.

We use the same symbol (e.g., A_circ^(1)) for both the constraint and its penalty term, following the σ-sector papers' notational convention; the distinction matters in §10 (Theorem 3), where the same expression $\oint_C \alpha$ plays opposite roles in different admissibility classes (constraint to be penalised in the dissipative class; conserved observable in the reversible class).

The on-sector master action under the dissipative class is therefore:

$$A_{cl}^{(1)}(\alpha) = \alpha^{(1)} (\oint_C \alpha)^2 + \gamma^{(1)} \langle \alpha, \Delta^1\alpha \rangle,$$

with coefficients $\alpha^{(1)}$, $\gamma^{(1)}$ determined by the substrate.

Comparing to the candidate functionals:

- **A_comp^(1)** ✓ admissible — corresponds to closure-competition (catalogue item).
- **A_circ^(1)** ✓ admissible — corresponds to circulation conservation (catalogue item).
- **A_anti^(1)** ✗ not in catalogue — no corresponding substrate constraint.
- **A_d2^(1)** ✗ not in catalogue at leading order — higher-derivative.
- **A_mag^(1)** ✗ not in catalogue — no substrate constraint penalising amplitude.

The absence of A_mag^(1) from the catalogue is the structural fact underlying Theorem 2. The $K = 7$ catalogue contains *structural* constraints (about how substrate variables relate to each other), not *amplitude* constraints — exactly as in the σ-sector papers' spoke variable catalogue, which contained no λ^2 mass-like term.

---

6. Lemma 1: Classification of Gradient-Flow Dynamics

Lemma 1 (Classification of gradient-flow dynamics — classical fact, used here to characterise the dissipative admissibility class). *Let $S[\phi]$ be a smooth action functional on a configuration space V, and let the dynamics be the gradient flow $\partial_\tau \phi = - \partial S[\phi] / \partial\phi$. Then for any initial configuration $\phi(0) \in V$, the evolution $\phi(\tau)$ is either:*

(i) ***invariant*** *($\phi(\tau) = \phi(0)$ for all $\tau \geq 0$), if $\phi(0)$ is a critical point of S; or*

(ii) ***strictly dissipative*** *($S[\phi(\tau)]$ is strictly decreasing for τ in some interval [0, T]), if $\phi(0)$ is not a critical point.*

Oscillatory dynamics is structurally excluded.

This is a classical fact about gradient flows. We state it as a Lemma rather than a Theorem because its content is standard — the substantive use is its application to the dissipative admissibility class, where it explains the structural shape of the σ-sector dynamics and, via the contrast with reversible-class dynamics, motivates the sector-classification.

Proof. Compute $dS/d\tau$ along the flow:

$$dS/d\tau = \langle \partial S/\partial\phi, \partial_\tau \phi \rangle = \langle \partial S/\partial\phi, - \partial S/\partial\phi \rangle = - \|\partial S/\partial\phi\|^2 \leq 0,$$

with equality iff $\partial S/\partial\phi = 0$ (critical point). If $\phi(0)$ is critical, $\partial_\tau \phi = 0$ and ϕ stays at $\phi(0)$ (case i). Otherwise $dS/d\tau < 0$ in a neighbourhood of $\tau = 0$, so S strictly decreases (case ii).

For oscillation: if $\phi(\tau)$ periodic with period $T > 0$, then $S[\phi(T)] = S[\phi(0)]$ and (since S is non-increasing) S is constant on [0, T], so $dS/d\tau = 0$ throughout, hence $\partial_\tau \phi = 0$ throughout, hence ϕ is constant — contradicting non-constant periodicity. □

Implication for the dissipative admissibility class. Combined with §4's identification of gradient flow as the natural variational principle of the dissipative admissibility class derived from the K = 7 master-action variation, Lemma 1 says that the dissipative class's dynamics is structurally limited to invariance and dissipation; oscillatory dynamics requires a different admissibility class. This is the structural motivation for the sector-classification: the persistent gauge sector's potential for non-trivial dynamics requires a different variational structure than gradient flow can provide, which Part II locates in the reversible admissibility class.

Application to the master dynamics on $C^1(C_6)$ under the dissipative class. Under the master action A_cl^(1) and gradient flow, the persistent gauge sector $\mathrm{Ker}(\delta^*)$ is either invariant or dissipative depending on whether A_cl^(1) has non-vanishing gradient on $\mathrm{Ker}(\delta^*)$.

On the harmonic subspace $\alpha = c\, \alpha_0$ (where $\alpha_0(e_i) = 1$):

- $\langle \alpha, \Delta^1\alpha \rangle = 0$ (Δ^1 vanishes on $\mathrm{Ker}(\Delta^1)$),
- $(\oint_C \alpha)^2 = 36c^2$,

so $A_{cl}^{(1)}(c\,\alpha_0) = 36\,\alpha^{(1)}\,c^2$. The gradient $\partial A_{cl}^{(1)}/\partial c = 72\,\alpha^{(1)}\,c$ is non-zero whenever $\alpha^{(1)} \neq 0$ and $c \neq 0$. Under gradient flow:

$$\partial_\tau c = -\,72\,\alpha^{(1)}\,c,$$

giving exponential decay $c(\tau) = c(0)\exp(-\,72\,\alpha^{(1)}\,\tau)$. The framework paper's invariance result corresponds to $\alpha^{(1)} = 0$ (omitting $A_{circ}^{(1)}$); with $\alpha^{(1)} > 0$, the harmonic mode is *dissipative* within the dissipative class.

In neither case is the dynamics oscillatory. The structural exclusion of oscillation by Lemma 1 is the dissipative class's defining feature — not a deficiency, but the appropriate domain of its variational principle. Different dynamics on $\mathrm{Ker}(\delta^*)$ belongs to a different admissibility class (Theorem 3, §10).

---

7. Theorem 2: The K = 7 Catalogue Does Not Generate Mass-Like Terms

Theorem 2 (The K = 7 catalogue does not generate mass-like terms within the dissipative class). *Mass-like quadratic terms $V(\alpha) = m^2\,\langle\alpha, \alpha\rangle$ penalising the magnitude of α — which do not vanish on Ker(δ)* — are not generated by the K = 7 constraint catalogue extended to 1-cochains. The catalogue admits a four-element closure ($A_{inc}^{(1)}$, $A_{hub}^{(1)}$, $A_{circ}^{(1)}$, $A_{comp}^{(1)}$) with no mass-like term.*

Important caveat: this theorem establishes only that the dissipative class's variational structure does not generate mass-like terms. It does not preclude Hamiltonian structure on the persistent sector, which arises from the reversible admissibility class via a different mechanism (the Wigner + Stone derivation of unitary evolution), not from K = 7 catalogue penalties.

Proof. The 1-cochain constraint catalogue (§5.3) closes at four constraints under (P5′). The two off-sector constraints ($A_{inc}^{(1)}$, $A_{hub}^{(1)}$) vanish on admissible 1-cochains. The two on-sector constraints are $A_{circ}^{(1)} = (\oint_C \alpha)^2$ and $A_{comp}^{(1)} = \langle\alpha, \Delta^1\alpha\rangle$.

A mass-like term $m^2\,\langle\alpha, \alpha\rangle$ corresponds to a substrate constraint of the form "α should be small everywhere," penalising amplitude at every edge. The K = 7 catalogue contains no such constraint: it contains *structural* constraints (relations between variables), not *amplitude* constraints. By (P5′), only catalogue items contribute to the admissible action within the dissipative class; therefore mass-like terms are not generated.

On the harmonic subspace:

- $A_{circ}^{(1)} = 36c^2$ (non-vanishing) — but this is a *constraint penalty*, not a *mass term*.
- $A_{comp}^{(1)} = 0$.

So the only catalogue-generated quadratic term not vanishing on $\mathrm{Ker}(\delta^*)$ is the circulation-conservation penalty, which functions as a constraint penalty rather than providing inertia. No positive-definite mass-like contribution $m^2 \langle\alpha, \alpha\rangle$ with $m^2 > 0$ is generated by the K = 7 catalogue within the dissipative class. □

Remark on what this theorem does and does not establish. The exclusion of mass-like terms from the K = 7 catalogue is the same structural feature that excluded them from the σ-sector papers' four-term decomposition of A_cl on the spoke variable. The catalogue is about *relations between variables*, not *amplitudes*. This is a load-bearing structural feature of the K = 7 architecture, consistent with the σ-sector's intrinsically massless envelope character.

The important refinement under the sector-classification framing: Theorem 2 establishes only that the dissipative class's variational structure (gradient flow on the K = 7 catalogue) does not generate mass-like terms. It does *not* preclude Hamiltonian structure on the persistent sector, because Hamiltonian structure belongs to the reversible admissibility class (Part II) and is generated by a different mechanism — the Wigner + Stone derivation establishing unitary evolution as the necessary form of reversible distinguishability-preserving evolution. The absence of mass-like terms from the dissipative catalogue is consistent with — not in contradiction with — non-trivial Hamiltonian dynamics on the persistent sector via the reversible admissibility class.

PART II — REVERSIBLE ADMISSIBILITY

8. The Reversible Admissibility Class

Reversible admissibility describes dynamics that preserve the substrate's distinguishability structure while transporting it through configuration space. The substrate moves through admissible configurations *without* dissipating structure: distinguishability content is preserved exactly, with no entropy reduction or convergence behaviour.

The characteristic features of reversible admissibility:

- **Distinguishability-preserving.** The inner-product structure on the (complexified) configuration space is preserved exactly under the dynamics.
- **Unitary evolution.** The time evolution is given by $U(\tilde{\tau}) = e^{-iH\tilde{\tau}}$ for a self-adjoint Hamiltonian operator H acting on the (complexified) configuration space.
- **Conservative.** The Hamiltonian H is conserved exactly along the flow; the dynamics has no Lyapunov functional decreasing toward equilibrium.
- **Recurrent.** Trajectories are bounded; for compact systems, recurrent rather than convergent.

These features are not optional structural choices; they are forced by the reversibility axioms themselves, as established by the Hamiltonian admissibility derivation summarised in §9 below.

The persistent gauge sector $\mathrm{Ker}(\delta^*) \cong H^1(C_6)$ of the master configuration space is the natural carrier of reversible admissibility within the K = 7 architecture. The Wilson-loop direction is topologically protected — it cannot be dissipated by any admissibility-restoring dynamics, because no substrate constraint penalises it (cf. Theorem 2). What remains is the question of what dynamics it *does* have, and the answer comes from the reversible admissibility class: unitary Hamiltonian evolution preserving the topological invariant.

---

9. The Hamiltonian Admissibility Derivation: Wigner + Stone

The Hamiltonian admissibility derivation (companion paper *The Hamiltonian as an Admissibility Derivation in VERSF*) establishes that unitary evolution $U(t) = e^{-iHt}$ with self-adjoint generator H is the *necessary* form of admissible reversible distinguishability-preserving evolution, derived as a representation-theorem consequence of three reversibility axioms. We summarise the derivation here as load-bearing background for Theorem 3 below; the full derivation is in the companion paper.

9.1 The three reversibility axioms

The derivation rests on three axioms applied to admissible reversible evolution of a system in a complex inner-product space $\mathscr{H}$ (the suitably complexified configuration space):

- **(A1) Composability.** Evolution through sequential time intervals composes consistently: $U(t_1 + t_2) = U(t_2) \circ U(t_1)$ for all $t_1, t_2 \geq 0$. This says that evolution is parametrised by an additive parameter (proper time, or its analogue in the relevant sector).
- **(A2) Continuity.** The map $t \mapsto U(t)$ is strongly continuous: $U(t)\,\psi \to \psi$ as $t \to 0$ for every $\psi \in \mathscr{H}$. This excludes discontinuous jumps in evolution.
- **(A3) Distinguishability preservation.** Inner products are preserved: $\langle U(t)\,\psi \mid U(t)\,\phi\rangle = \langle\psi \mid \phi\rangle$ for all $\psi, \phi \in \mathscr{H}$ and all $t \geq 0$. This is the formal expression of distinguishability preservation: states that are distinguishable (orthogonal) at one time remain distinguishable; states with given overlap retain that overlap.

9.2 The derivation

The three axioms force the dynamical structure through the following chain:

Step 1 (Wigner's theorem). Distinguishability preservation (A3) at each fixed t means each U(t) is a symmetry of the inner-product structure on $\mathscr{H}$. By Wigner's theorem, any such symmetry is either unitary or antiunitary (up to a phase). So U(t) for each t is either unitary or antiunitary.

Step 2 (Continuity excludes antiunitary). By composition (A1), $U(t) = U(t/2) \circ U(t/2)$. The composition of two antiunitary maps is unitary (the antilinear conjugations cancel), so if U(t/2) is antiunitary, U(t) is unitary. By continuity (A2), $U(t) \to I$ (identity, unitary) as $t \to 0$. A path from antiunitary to unitary through the $U(t/2) \circ U(t/2)$ decomposition cannot remain continuous if U(t) takes both unitary and antiunitary values, because the unitary and antiunitary operators form disconnected components in the operator topology. So U(t) is unitary for all t (the antiunitary branch is excluded).

Step 3 (Stone's theorem). Composability (A1) means t ↦ U(t) is a one-parameter group of unitary operators on $\mathscr{H}$. Strong continuity (A2) makes it a strongly continuous one-parameter unitary group. By Stone's theorem, every such group has a unique self-adjoint generator H such that:

U(t) = e^(−iHt) for all t ∈ ℝ,

with the sign convention chosen so that H has spectrum bounded below for physical systems.

Step 4 (The Hamiltonian). The generator H is, by definition, the *Hamiltonian* of the system. The Schrödinger equation i*ℏ* ∂_t ψ = H ψ emerges as the differential form of U(t) ψ.

9.3 Summary

The chain of derivation:

Composability + Continuity + Distinguishability preservation → Wigner: each U(t) unitary or antiunitary → Continuity: antiunitary excluded, U(t) unitary → Stone: U(t) = e^(−iHt) with unique self-adjoint H → Schrödinger evolution iℏ ∂_t ψ = H ψ.

The Hamiltonian is therefore not an independent postulate but a *representation-theorem consequence* of the three reversibility axioms. Within the reversible admissibility class, unitary Hamiltonian evolution is forced — it is the unique mathematical structure consistent with reversible distinguishability-preserving evolution.

This is the load-bearing content imported from the Hamiltonian admissibility derivation. Within the present paper, we take this derivation as established and apply it to the persistent gauge sector in Theorem 3.

---

10. Theorem 3: Reversible Admissibility on the Persistent Sector

Theorem 3 (Reversible admissibility structure on Ker(δ)). The persistent gauge sector Ker(δ) ≅* $H^1(C_6)$ ≅ ℝ is the natural carrier of reversible admissibility within the master configuration space $C^1(C_6)$, with the assignment to the reversible admissibility class treated as the structural postulate flagged in §3. Under the Hamiltonian admissibility derivation (§9) — which establishes U($\tilde{\tau}$) = e^(−iH$\tilde{\tau}$) with self-adjoint H as the necessary form of reversible distinguishability-preserving evolution — the persistent direction admits Hamiltonian generator structure of the form:*

α_harm($\tilde{\tau}$) = U($\tilde{\tau}$) α_harm(0) where U($\tilde{\tau}$) = e^(−iH_h $\tilde{\tau}$),

for a self-adjoint operator H_h acting on Ker(δ) (suitably complexified to a complex inner-product space). The dynamics:*

- *preserves the Wilson-loop topological invariant $\oint_C \alpha_{harm}$ exactly as the conserved scalar observable;*
- *is non-dissipative: distinguishability content on Ker(δ)* is preserved exactly;*
- *takes the explicit form $U(\tilde{\tau}) = e^{-i\omega\tilde{\tau}}$ with ω a real scalar (the most general self-adjoint operator on a one-dimensional complex space). The case ω > 0 gives recurrent oscillation with period 2π/ω; the case ω = 0 reduces to trivially constant evolution, recovering the framework paper's invariance result as a perfectly valid special case.*

The specific value of ω — including the question of whether ω = 0 or ω > 0 — is undetermined by the reversibility axioms (A1)–(A3) alone on the bare one-dimensional space. Theorem 3 therefore establishes the form of the generator structure (unitary, with single-scalar self-adjoint generator) but does not pin down its eigenvalue. Additional structural input — from the persistent cohomological transport identification of Part III, or from inter-class coupling in the broader architecture (R3 in Theorem 4, §16), or from matter coupling — is required to fix ω.

Proof. The Hamiltonian admissibility derivation (§9) establishes that reversible distinguishability-preserving evolution on any (suitably complexified) inner-product configuration space takes the form $U(\tilde{\tau}) = e^{-iH\tilde{\tau}}$ for some self-adjoint H.

Applied to the persistent gauge sector: $\mathrm{Ker}(\delta^*) \cong H^1(C_6) \cong \mathbb{R}$ is one-dimensional. The complexified persistent sector $\mathscr{H}_h = \mathrm{Ker}(\delta^*) \otimes \mathbb{C} \cong \mathbb{C}$ admits a natural inner product inherited from the cochain inner product on $C^1(C_6)$ (the L^2 inner product on edge functions). Reversible admissibility on $\mathscr{H}_h$ satisfies axioms (A1)–(A3), so by the Hamiltonian admissibility derivation, the evolution takes the form $U(\tilde{\tau}) = e^{-iH_h \tilde{\tau}}$ for self-adjoint H_h on $\mathscr{H}_h$.

Since $\mathscr{H}_h$ is one-dimensional (over $\mathbb{C}$), the most general self-adjoint H_h is multiplication by a real scalar ω. So:

$$U(\tilde{\tau})\, \alpha_{harm} = e^{-i\omega\tilde{\tau}}\, \alpha_{harm}.$$

Conservation of the Wilson loop: $|U(\tilde{\tau})\, \alpha_{harm}|^2 = |\alpha_{harm}|^2$, so the magnitude of the harmonic mode (which is the Wilson loop, up to normalisation) is conserved exactly. Distinguishability preservation: $\langle U(\tilde{\tau})\, \alpha \mid U(\tilde{\tau})\, \beta\rangle = \langle \alpha \mid \beta\rangle$ for $\alpha, \beta \in \mathscr{H}_h$. Recurrence (for ω > 0): $U(\tilde{\tau} + 2\pi/\omega) = U(\tilde{\tau})$, so periodic with period 2π/ω. For ω = 0, $U(\tilde{\tau})$ = identity (trivially constant, recovering the framework paper's invariance as a special case). □

Discussion. Theorem 3 establishes the *form* of the Hamiltonian generator structure on the persistent sector, derived from a structurally distinct source than the K = 7 catalogue: the reversible admissibility class's variational structure, as captured by the Hamiltonian admissibility derivation. What is *not* established by Theorem 3 alone is the specific eigenvalue ω of the generator. Both ω = 0 (giving the framework paper's invariance result as a special case) and ω > 0 (giving recurrent oscillation) are permitted by the reversibility axioms on the bare 1D persistent sector; nothing in Theorem 3 forces one or the other.

The contrast with the framework paper: the framework paper applied the gradient-flow variational structure uniformly across the master configuration space, getting trivial invariance on Ker(δ*) as the necessary consequence of that uniform postulate. Theorem 3 makes the alternative postulate (Hamiltonian generator structure on Ker(δ*) from the reversible admissibility class) substantive by sourcing it from the Hamiltonian admissibility derivation. The framework paper's invariance result is now seen as the $\omega = 0$ special case of Theorem 3 — permitted by the alternative postulate, not excluded by it.

The contrast with Theorem 2: Theorem 2 established that the $K = 7$ catalogue (a feature of the dissipative class) does not generate mass-like terms. Theorem 3 establishes that *the form* of Hamiltonian generator structure on the persistent sector is well-defined despite this absence, because the form arises from reversibility (a different mechanism in a different admissibility class), not from $K = 7$ catalogue penalties. The two theorems are consistent and complementary, not contradictory.

The structural-permits-not-forces character of Theorem 3 is the same kind of result as the framework paper's kinematic-not-dynamical character: a real architectural contribution, with the specific quantitative content (ω here, ω-equivalent there) punted to subsequent work. Where Part III enters: by identifying Ker(δ*) as the $K = 7$ specialisation of the broader substrate's refinement-stable observable sector $H^1(G(\Lambda))$, and applying the Strong Synthesis Theorem of the Maxwell admissibility paper, we obtain a substantively richer characterisation of the persistent sector's dynamics in the continuum limit — specifically, that the continuum-limit form is Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$. This addresses the *form* of the continuum dynamics; specific frequencies and coupling constants remain dependent on substrate-scale-to-continuum-scale mapping and matter coupling.

Remark on the structural inversion of ∮_C α between the two admissibility classes. A striking structural feature of the two-class architecture: the same expression ∮_C α plays *opposite* roles in the dissipative and reversible admissibility classes.

- In the *dissipative class*, ∮_C α appears squared as the constraint-violation penalty A_circ^(1) = $(\oint_C \alpha)^2$ in the master action A_cl^(1). Under gradient flow, this penalty drives ∮_C α toward zero — i.e., it pushes the configuration *away from* the harmonic generator α_0 (which has $\oint_C \alpha_0 = 6$) toward configurations with zero circulation (which are in the orthogonal exact subspace Im(δ)). The Wilson loop is *penalised*.
- In the *reversible class*, ∮_C α appears as the *conserved scalar observable* of the Hamiltonian dynamics (Theorem 3, preservation of the Wilson loop). Unitary evolution preserves $|\alpha\_harm|^2$ and hence ∮_C α_harm exactly. The Wilson loop is *conserved*.

The same expression — Lemma 3.1's circulation map from the framework paper — therefore plays a constraint-penalty role in one admissibility class and a conserved-observable role in the other. This is not a contradiction but a structural feature of the two-class architecture: the dissipative class operates on the full configuration space $C^1(C_6)$, with the harmonic subspace being a "target to be moved away from" via the circulation penalty; the reversible class operates on the harmonic subspace itself, with the Wilson loop being its conserved observable. The two classes look at the same expression from opposite sides of the orthogonal decomposition Im(δ) ⊕ Ker(δ*).

PART III — PERSISTENT COHOMOLOGICAL TRANSPORT AND MAXWELL-FORM U(1)

11. The Persistent Cohomological Sector and Wilson/Cohomology Identification

The framework paper's identification of the persistent gauge sector with the first cohomology $H^1(C_6)$ of the rim cycle is the K = 7 specialisation of a more general identification developed elsewhere in the VERSF programme: that the persistent observable sector of any admissible substrate is the first cohomology $H^1(G(\Lambda))$ of the substrate's underlying admissibility complex $G(\Lambda)$. This section makes the K = 7 → general substrate identification explicit, since it is the bridge that allows the Maxwell admissibility theorem (which operates on $H^1(G(\Lambda))$) to apply to the σ-sector framework's $\mathrm{Ker}(\delta^*)$.

11.1 The substrate-level persistent sector

The broader VERSF programme establishes (in the substrate-level papers, particularly those on refinement persistence and the Fold and the Record architecture) that for any admissible substrate Λ with admissibility complex $G(\Lambda)$:

- The first cohomology $H^1(G(\Lambda))$ is the unique refinement-stable observable sector — meaning, the unique structure that survives the refinement operator Γ^* unchanged (up to the structural scaling factor: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$, as established in those papers).
- Cochain representatives $\omega \in C^1(G(\Lambda))$ of cohomology classes $[\omega] \in H^1(G(\Lambda))$ play the role of the *potential* in the conventional gauge-theoretic sense, with the equivalence relation $\omega \sim \omega + d^0\varphi$ (gauge equivalence) forced by the trivialisation of the scalar sector C^0.
- The natural physical observable is the holonomy / Wilson-loop integral $\oint_\gamma \omega \cdot d\ell$ around closed homology cycles $\gamma \in H_1(G(\Lambda))$. By the homology-cohomology pairing, this is a topological invariant depending only on the classes $[\gamma]$ and $[\omega]$.

11.2 The K = 7 specialisation

The σ-sector framework's master configuration space $C^1(C_6)$ is the K = 7 specialisation of this structure, with $G(\Lambda\_K{=}7)$ being the K = 7 architecture's rim cycle C_6:

- The persistent gauge sector subspace $\mathrm{Ker}(\delta^*) \subset C^1(C_6)$ is the harmonic representative subspace of $H^1(C_6) \cong \mathbb{R}$, by the Hodge theorem on cochain complexes.
- The harmonic 1-cochain α_0 (with $\alpha_0(e_i) = 1$) is the natural representative of the unique non-trivial cohomology class $[\alpha_0] \in H^1(C_6)$.
- The Wilson-loop integral $\oint_C \alpha = \Sigma_i\, \alpha(e_i)$ is the conserved scalar observable. By the framework paper's Lemma 3.1, it is well-defined on cohomology classes (invariant under $\alpha \to \alpha + \delta\chi$ for $\chi \in C^0(C_6)$) and gives the isomorphism $H^1(C_6) \cong \mathbb{R}$.

The K = 7 framework is therefore the discrete K = 7 special case of the more general substrate framework. The persistent gauge sector Ker(δ*) of the K = 7 framework is the K = 7 specialisation of $H^1(G(\Lambda))$.

11.3 The three input identifications, plus the Maxwell-form result

The integration between the Maxwell admissibility paper and the substrate-level work has been carried out in a separate paper (the *Strong Synthesis paper*, cited in §13). It establishes three structural identifications — (I)–(III) below — that serve as *inputs* to the Strong Synthesis Theorem, together with the *output* result (IV) that the Strong Synthesis Theorem establishes:

Input identifications (the bridge):

- **(I) The potential is the cochain representative.** The Maxwell paper's transport potential A_c^μ is the cochain representative of $[\omega] \in H^1(G(\Lambda))$ under the IR limit identifying continuum 1-forms with substrate edge cochains.
- **(II) Gauge redundancy is refinement equivalence.** The Maxwell paper's gauge transformation A_c → A_c + ∂χ is the equivalence relation $\omega \sim \omega + d^0\varphi$ on cochains.
- **(III) Wilson loops are persistent observables.** The closed-walk integral ∮_γ A_c · dℓ is the substrate-level pairing $\langle[\gamma], [\omega]\rangle : H\_1(G(\Lambda)) \times H^1(G(\Lambda)) \to \mathbb{R}$. Refinement persistence makes Wilson loops refinement-stable up to the structural scaling factor.

Output result (established by the Strong Synthesis Theorem of §13, not by the identifications themselves):

- **(IV) Maxwell-form dynamics on the persistent sector** — see the Strong Synthesis Theorem in §13.

The three input identifications (I)–(III) are the bridge that makes the Strong Synthesis composition possible: they show that the Maxwell admissibility paper's "transport potential" and the σ-sector framework's "persistent gauge sector cochain" are the same object under two complementary descriptions, that the gauge-redundancy condition (B2) of the Maxwell admissibility theorem coincides with the substrate-level refinement equivalence, and that Wilson loops are the natural conserved observables of both sides. The output result (IV) is what the Strong Synthesis Theorem establishes by composing the Maxwell admissibility theorem's uniqueness with the substrate-level uniqueness via (I)–(III).

For the K = 7 special case, the identifications become:

- The 1-cochain α on rim edges is the cochain representative of $[\alpha] \in H^1(C_6)$.
- The gauge equivalence $\alpha \sim \alpha + \delta\chi$ for $\chi \in C^0(C_6)$ is the K = 7 specialisation of $\omega \sim \omega + d^0\varphi$.
- The Wilson loop $\oint\_C \alpha = \Sigma_i \alpha(e_i)$ is the K = 7 specialisation of ∮_γ A_c · dℓ for γ the rim cycle.
- Maxwell-form dynamics on Ker(δ*) is the K = 7 specialisation of Maxwell-form U(1) transport on $H^1(G(\Lambda))$ — to be established in §14.

12. The Maxwell Admissibility Uniqueness Theorem

The Maxwell admissibility paper (cited throughout as the *Maxwell admissibility paper*) establishes that under the substrate axioms BCB and TPB, within a specific admissibility class motivated by substrate-level informational structure, Maxwell-form U(1) gauge transport is the unique surviving continuum dynamics. We summarise the theorem here as load-bearing background for the Main Synthesis Theorem of §14.

12.1 The substrate axioms

- **BCB (Bit Conservation and Balance).** Committed distinctions are locally conserved. They may flow, redistribute, or reorganise, but cannot be created or annihilated arbitrarily. Differentially: $\partial_t s + \nabla\cdot J_s = 0$ for an entropy / commitment density s and current J_s.
- **TPB (Ticks-Per-Bit).** The propagation of distinguishability is mediated by substrate update progression with a finite minimum cost per resolved bit, implying a universal finite-speed bound on commitment transport. In the continuum limit, this becomes a Lorentzian causal structure with a maximum propagation speed.

These two axioms are the *substrate-level* admissibility axioms of the broader VERSF programme. They are derived from finite distinguishability + irreversible commitment + finite-cost-per-resolved-bit, as established in the substrate-level papers.

12.2 The admissibility class (B1)–(B4)

The Maxwell admissibility paper develops an admissibility class for transport dynamics on the persistent observable sector under BCB/TPB, with four constraints:

- **(B1) Locality.** Transport dynamics depends on local values of fields and their derivatives at the substrate scale, not on arbitrary non-local integrals. This is the continuum-limit version of (P1b) from the σ-sector papers.
- **(B2) Gauge redundancy.** The dynamics respects the gauge equivalence $A_c \to A_c + \partial\chi$ inherited from the persistent sector's cochain structure (cf. §11). This is the continuum-limit version of the refinement-equivalence on $H^1(G(\Lambda))$.
- **(B3) Closure-geometry covariance.** The dynamics respects the substrate's closure geometry (the geometric structure of admissible configurations), inheriting Lorentz covariance from the TPB causal bound in the continuum limit.
- **(B4) ξ-EFT truncation.** The dynamics is the leading-order term in a controlled effective-field-theory expansion in a small parameter ξ (related to the substrate-scale-to-continuum-scale ratio), with the leading-order being $O(\varepsilon^0)$ where $\varepsilon = \xi$ is the lattice spacing analogue.

Status of (B1)–(B4) vis-à-vis BCB and TPB. A point that affects the strength of the Main Synthesis Theorem deserves explicit attention: the four admissibility-class conditions are *substrate-motivated by* BCB and TPB but are not strictly *derived from* them. They are independent admissibility postulates layered on top of the substrate axioms, with the substrate axioms providing motivation rather than logical derivation. The Maxwell admissibility paper is explicit about this distinction:

- **(B1) Locality** is motivated by the substrate's lattice structure (substrate constraints are local to admissibility cells) but not derived: it is conceivable that emergent continuum dynamics could carry non-local structure even on a local substrate, and (B1) is the postulate that this does not happen at leading order. It is a natural postulate but not forced.
- **(B2) Gauge redundancy** is motivated by — and arguably approximately derived from — the substrate-level refinement equivalence $\omega \sim \omega + d^0\varphi$ on cochains (which is forced by the scalar trivialisation in the substrate-level papers). The §11 four identifications establish this linkage at the level of structure; whether the substrate-level refinement equivalence forces continuum gauge invariance in the strict sense of (B2), or only motivates it as a natural continuum-limit translation, is a substantive technical question that the Strong Synthesis paper addresses but does not fully close.
- **(B3) Closure-geometry covariance** is motivated by TPB (which imposes a finite-speed bound on commitment transport, giving the continuum limit a Lorentzian causal structure) but the continuum Lorentz covariance is itself a strong assumption layered on top of TPB. The substrate-level causal bound translates to "no faster-than-c propagation" in the continuum, but full Lorentz covariance of the dynamics (not just kinematics) is an additional structural input.
- **(B4) ξ-EFT truncation** is a standard EFT methodological postulate, not specific to VERSF. It is justified by the substrate-to-continuum scale ratio being small, but the truncation to $O(\varepsilon^0)$ is a methodological choice.

The honest assessment: BCB and TPB *motivate* (B1)–(B4) and provide substrate-level grounding for them, but the four conditions are *not derived* from BCB/TPB in the strict sense. The Main Synthesis Theorem's force therefore depends on (B1)–(B4) being natural admissibility postulates well-motivated by the substrate axioms, rather than on them being logically forced by the substrate axioms. This is a substantive but bounded strength: stronger than "Maxwell-form is the unique standard-EFT theory for a U(1) gauge field on Lorentzian background" (because the admissibility conditions have substrate-level motivation rather than just renormalisability or experimental motivation), but weaker than "Maxwell-form is logically forced by BCB/TPB alone" (which would require deriving (B1)–(B4) from the substrate axioms, which is not done).

Whether (B1)–(B4) can be tightened — i.e., derived more strictly from BCB/TPB rather than merely motivated — is an open structural question in the broader programme. The present paper takes the Maxwell admissibility paper's stated conditions at face value and propagates the honest assessment of their status into the Main Synthesis Theorem's caveats.

12.3 The Maxwell admissibility theorem

Theorem (Maxwell admissibility uniqueness — cited result, Maxwell admissibility paper). *Under substrate axioms BCB and TPB, within the admissibility class (B1)–(B4), the unique* *surviving transport theory at $O(\varepsilon^0)$ on the persistent observable sector is Maxwell electrodynamics with U(1) gauge structure:*

$$\partial_\mu F^{\mu\nu} = J^\nu, \quad F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu,$$

where A^μ is the persistent-sector potential (the continuum limit of the cochain representative of $H^1(G(\Lambda))$), $F^{\mu\nu}$ is the field strength, and J^ν is the conserved commitment current (the continuum limit of the substrate's local distinguishability current).

The theorem establishes Maxwell-form *uniqueness*: no other admissible transport theory at $O(\varepsilon^0)$ is consistent with BCB, TPB, and (B1)–(B4). The result does not assume Maxwell's equations; it derives them as the unique closure-compatible transport geometry under the admissibility constraints.

The proof in the companion paper proceeds via Helmholtz decomposition of the commitment current under BCB, identification of the field tensor's antisymmetric structure from (B2) gauge redundancy and (B3) closure-geometry covariance, and (B4) ξ-EFT truncation excluding higher-derivative or non-linear corrections at leading order. The Lorentz covariance is inherited from TPB; the U(1) structure is inherited from the abelian gauge redundancy (B2), which itself is the continuum limit of the refinement equivalence on H^1.

---

13. The Strong Synthesis Theorem: Maxwell-Form on $H^1(G(\Lambda))$

A separate paper in the broader VERSF programme — the substrate–Maxwell integration paper, which we cite as the *Strong Synthesis paper* — integrates the Maxwell admissibility uniqueness theorem (§12) with the substrate-level identification of $H^1(G(\Lambda))$ as the unique refinement-stable observable sector (§11), yielding the Strong Synthesis Theorem. We summarise the theorem here together with a proof sketch showing how the three input identifications of §11.3 mediate the composition; the full derivation is in the Strong Synthesis paper.

Theorem (Strong Synthesis — cited result, Strong Synthesis paper). *Under substrate axioms BCB and TPB, within the admissibility class (B1)–(B4), Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$ is the unique refinement-stable continuum-limit dynamics on the unique refinement-stable observable sector $H^1(G(\Lambda))$.*

Proof sketch (composition via the three input identifications of §11.3). The Maxwell admissibility theorem (§12) gives uniqueness of the *dynamics* (Maxwell-form U(1) at $O(\varepsilon^0)$) given a carrier of "the right type" — specifically, a 1-form gauge field A^μ on a Lorentzian substrate-continuum manifold, subject to gauge equivalence $A^\mu \to A^\mu + \partial^\mu \chi$. The substrate-level work gives uniqueness of the *carrier* ($H^1(G(\Lambda))$) given the scalar trivialisation that forces refinement equivalence on cochains: $\omega \sim \omega + d^0\varphi$. The combination requires showing that these two specifications describe the same object — that the "right type of carrier" assumed by the Maxwell admissibility theorem and the refinement-stable observable sector identified by the substrate-level work are not merely analogous but are literally the same structure in the IR continuum limit. The three input identifications of §11.3 establish exactly this:

- **(I) The potential is the cochain representative.** The Maxwell admissibility theorem's potential A^μ is, under the IR limit identifying continuum 1-forms with substrate edge cochains, exactly the cochain representative of the cohomology class $[\omega] \in H^1(G(\Lambda))$. The "1-form gauge field" of the Maxwell admissibility theorem and the "cochain representative of the persistent sector" of the substrate-level work are the same object. This is the *carrier identification*.
- **(II) Gauge redundancy is refinement equivalence.** The Maxwell admissibility theorem's (B2) condition — gauge equivalence $A^\mu \to A^\mu + \partial^\mu \chi$ — is, under the cochain-to-1-form identification of (I), exactly the refinement equivalence $\omega \sim \omega + d^0\varphi$ on cochains. The "gauge redundancy" of the Maxwell admissibility theorem and the "refinement equivalence" of the substrate-level work are the same equivalence relation. This is the *gauge identification*.
- **(III) Wilson loops are persistent observables.** The Maxwell admissibility theorem's gauge-invariant observable $\oint_\gamma A^\mu dx_\mu$ is, under the identifications of (I)–(II), exactly the homology-cohomology pairing $\langle[\gamma], [\omega]\rangle : H_1(G(\Lambda)) \times H^1(G(\Lambda)) \to \mathbb{R}$ of the substrate-level work. Refinement persistence ($\Gamma^* \circ L = 2\cdot\mathrm{Id}$, from the substrate-level papers) makes Wilson loops refinement-stable up to the structural scaling factor. The "gauge-invariant observable" of the Maxwell admissibility theorem and the "persistent observable" of the substrate-level work are the same physical quantity. This is the *observable identification*.
- **(IV) Composition.** Under (I)–(III), the carrier assumed by the Maxwell admissibility theorem and the carrier identified by the substrate-level work coincide. The Maxwell admissibility theorem's uniqueness of the dynamics on its assumed carrier therefore directly transfers to uniqueness of the dynamics on the substrate-level work's identified carrier: Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$ is the unique admissible continuum-limit dynamics on $H^1(G(\Lambda))$. □

The Strong Synthesis Theorem is substantially stronger than either input theorem read alone. The Maxwell admissibility theorem proves uniqueness of the *dynamics* given a carrier of the right type. The substrate-level papers prove uniqueness of the *carrier* given the scalar trivialisation. The Strong Synthesis combines them: under the four §11.3 identifications, *both* the carrier ($H^1(G(\Lambda))$) *and* the dynamics (Maxwell-form U(1)) are unique under the substrate axioms and admissibility class, with the bridging identifications themselves derivable from the BCB/TPB-motivated structure of admissible refinement.

Conditions of validity (explicit caveats). The Strong Synthesis Theorem inherits its conditions from its three component pieces: (i) the substrate axioms BCB and TPB, (ii) the admissibility class (B1)–(B4) of the Maxwell admissibility theorem (with §12.2 caveats about which conditions are derived from BCB/TPB and which are independent postulates), (iii) the IR continuum-limit identification of substrate edge cochains with continuum 1-forms (which itself requires the substrate-to-continuum bridge of the substrate-level papers). Readers relying on the Main Synthesis Theorem of §14 should consult the Strong Synthesis paper for the full derivation, particularly for the technical details of the IR limit and the substrate-to-continuum bridge.

The Strong Synthesis Theorem is the load-bearing positive content of Part III. The Main Synthesis Theorem of §14 specialises it to the K = 7 framework.

---

14. Main Synthesis Theorem: Maxwell-Form Persistent Transport on the σ-Sector Framework

Main Synthesis Theorem (Maxwell-form persistent transport on the σ-sector framework). *The persistent gauge sector Ker(δ)* $\cong H^1(C_6)$ of the σ-sector master-action framework, considered as the K = 7 specialisation of the general substrate's refinement-stable observable sector $H^1(G(\Lambda))$:*

(i) [Reversible admissibility class membership] supports admissible reversible distinguishability-preserving evolution, as established by Theorem 3 of this paper.

(ii) [Hamiltonian generator necessity] necessarily carries unitary Hamiltonian transport $U(\tilde{\tau}) = e^{(-iH\_h \tilde{\tau})}$ generated by a self-adjoint operator H_h, as established by the Hamiltonian admissibility derivation (Wigner + Stone, §9).

(iii) [Maxwell-form uniqueness in continuum limit] under substrate axioms BCB and TPB, within the admissibility class (B1)–(B4), the unique refinement-stable continuum-limit reversible dynamics on the persistent sector is Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$, as established by the Strong Synthesis Theorem (§13).

Composing (i), (ii), (iii): under BCB, TPB, and (B1)–(B4), the σ-sector master-action framework's persistent gauge sector — considered in its continuum limit as a refinement-stable observable sector — carries Maxwell-form U(1) gauge transport as the unique admissible reversible continuum dynamics.

Proof. The theorem is established by composition of three results, two of which are derived in companion papers and one (Theorem 3) within the present paper.

Step 1 — Reversible admissibility class membership (Theorem 3). The persistent gauge sector $Ker(\delta^*) \subset C^1(C_6)$ of the framework paper is identified as the K = 7 specialisation of $H^1(C_6)$ via the Hodge theorem (framework paper §5.3). Under the sector-classification framing of this paper (§3 dissipative class, §8 reversible class), Ker(δ*) is the natural carrier of the reversible admissibility class within the master configuration space (Theorem 3, §10).

Step 2 — Hamiltonian generator necessity (Hamiltonian admissibility derivation, §9). Under axioms (A1) composability, (A2) continuity, and (A3) distinguishability preservation applied to the suitably complexified Ker(δ*), the Hamiltonian admissibility derivation (Wigner + Stone) establishes that the dynamics takes unitary form $U(\tilde{\tau}) = e^{(-iH\_h \tilde{\tau})}$ with self-adjoint H_h (Theorem 3).

Step 3 — Identification with $H^1(G(\Lambda))$ (§11). The K = 7 framework's persistent gauge sector $\mathrm{Ker}(\delta^*) \cong H^1(C_6)$ is the K = 7 specialisation of the broader substrate's refinement-stable observable sector $H^1(G(\Lambda))$, via the three structural identifications of §11.3 that serve as *inputs* to the Strong Synthesis Theorem: (I) the potential is the cochain representative, (II) gauge redundancy is refinement equivalence, (III) Wilson loops are persistent observables. (Note: §11.3 lists a fourth item — "Maxwell-form dynamics on the persistent sector" — which is the *output* of the Strong Synthesis Theorem rather than an input identification; we list only the three input identifications here.)

Step 4 — Maxwell-form uniqueness (Strong Synthesis Theorem, §13). Under substrate axioms BCB and TPB, within the admissibility class (B1)–(B4), the Strong Synthesis Theorem establishes that the unique refinement-stable continuum-limit dynamics on $H^1(G(\Lambda))$ is Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$.

Composition. Steps 1–2 establish that the K = 7 framework's persistent gauge sector carries reversible admissibility dynamics in the form of unitary Hamiltonian evolution. Step 3 identifies this sector as the K = 7 specialisation of $H^1(G(\Lambda))$. Step 4, applied to this identification, establishes that the unique admissible continuum-limit reversible dynamics is Maxwell-form U(1) gauge transport.

Combining: under BCB, TPB, and (B1)–(B4), the σ-sector framework's persistent gauge sector carries Maxwell-form U(1) gauge transport as the unique admissible reversible continuum dynamics on the persistent observable sector. □

What "Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$" means on the K = 7 specialisation. A careful reader will notice an issue requiring explicit treatment: the Strong Synthesis Theorem operates on $H^1(G(\Lambda))$ for a *general* substrate, where H^1 can have higher rank, the persistent sector carries multiple modes, and the field equations have spatial-derivative content (giving wave propagation, Maxwell's equations in their full $\partial_\mu F^{\mu\nu} = 0$ form, etc.). The K = 7 specialisation collapses much of this content: $H^1(C_6) \cong \mathbb{R}$ is one-dimensional, the persistent sector $\mathrm{Ker}(\delta^*)$ carries a single mode (the harmonic generator α_0), and there is no spatial-derivative structure to support wave propagation. What survives the K = 7 specialisation?

What survives:

- **The kinematic structure of the persistent gauge sector as the natural carrier of U(1)-gauge-theoretic content** — the cochain-representative-as-potential, the cohomology-class-as-gauge-equivalence-orbit, the Wilson-loop-as-conserved-observable. These are kinematically present in the K = 7 case (with one cohomology class instead of many).
- **The structural identification of the reversible dynamics on the persistent sector as belonging to the same admissibility class that the Strong Synthesis Theorem characterises Maxwell-form on $H^1(G(\Lambda))$** — i.e., the K = 7 specialisation is the 1D specialisation of the same architectural setup, not a different kind of structure.
- *The Hamiltonian generator structure $U(\tilde{\tau}) = e^{-i\omega\tilde{\tau}}$ on $\mathrm{Ker}(\delta)$ from Theorem 3** — which is the 1D specialisation of the Hamiltonian generator that, in higher dimensions, would generate the full Maxwell wave-equation evolution on H^1.

What does *not* survive the $K = 7$ specialisation:

- **Substantive wave propagation content.** The free Maxwell equations $\partial_\mu F^{\mu\nu} = 0$ require spatial-derivative content; on a 1D persistent sector with one harmonic mode, there is no spatial structure for waves to propagate in.
- **Higher-rank gauge field content.** Conventional Maxwell theory has the gauge field A^μ as a 4-vector with spacetime indices; the $K = 7$ specialisation has only one cochain coefficient.
- **Field-strength-tensor content.** $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$ is non-trivial only with multiple spacetime indices; the 1D specialisation collapses this to a single scalar.

The honest interpretation of "Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$" on the $K = 7$ specialisation is therefore: **the structural template of Maxwell-form U(1) gauge theory restricted to its 1D harmonic-mode specialisation**, with the substantive wave-propagation content awaiting higher-dimensional generalisations (P3 in §22). The Main Synthesis Theorem's full Maxwell-form content lives in the higher-dimensional version (general substrate with rich $H^1(G(\Lambda))$); the $K = 7$ specialisation gives the structural template within which that content fits, but does not itself instantiate the propagating wave dynamics that Maxwell-form denotes in its full sense.

This is a genuine restriction of the Main Synthesis Theorem's $K = 7$ content, and we flag it explicitly. The substantive payoff of the theorem is structural — identifying that the $K = 7$ framework's persistent sector belongs to the same admissibility-class architecture that, in higher dimensions, gives the full Maxwell theory — not the immediate instantiation of conventional electromagnetism on the $K = 7$ rim. The latter requires the higher-dimensional generalisations (P3, P10) that are flagged as open problems.

Discussion of the composition. The composition is non-trivial. Each of the three load-bearing results in the composition (Theorem 3, the Hamiltonian admissibility derivation, the Strong Synthesis Theorem) is established separately in its respective setting; what makes the composition work — and what constitutes the substantive content of the present paper — is three structural contributions that link them:

- The *sector-classification* (Part I) recognising that the persistent gauge sector belongs to the reversible admissibility class, distinct from the dissipative class that governs the σ-sector. This is the substantive new conceptual content of this paper.
- The *Hamiltonian admissibility derivation applied to the persistent sector specifically* (Theorem 3, §10). This is the application of the companion derivation to the specific configuration space $\mathrm{Ker}(\delta^*)$ of the framework, with the assignment treated as the structural postulate flagged in §3.
- The *identification of* $\mathit{Ker}(\delta)$ as the $K = 7$ specialisation of $H^1(G(\Lambda))^*$ (§11). This bridges the $K = 7$ framework to the broader substrate framework on which the Maxwell admissibility theorem operates, via the three input identifications of §11.3.

These three contributions are the paper's substantive content: they are what makes the composition possible. The Main Synthesis Theorem itself is then the composition of these three contributions with the three load-bearing cited results, giving the positive structural result on the persistent sector.

Honest scope of the Main Synthesis Theorem. The theorem establishes Maxwell-form U(1) gauge transport on the persistent sector *in the continuum limit*, at $O(\varepsilon^0)$, under BCB/TPB and (B1)–(B4), with the K = 7 specialisation being the 1D restriction of this Maxwell-form structure (as detailed above). It does NOT establish:

- *Substantive wave-propagation content on the K = 7 framework itself.* Per the discussion above, this requires higher-dimensional generalisations.
- *Full relativistic completion.* The Maxwell admissibility theorem gives Maxwell-form at $O(\varepsilon^0)$; higher-order corrections, non-Abelian extensions, and full relativistic field theory beyond Maxwell are not addressed.
- *Matter coupling.* The theorem gives the pure-gauge-sector dynamics. Coupling to charged matter (electrons, sources, currents from non-substrate-level matter sectors) requires separate derivation — this is the matter-coupling question of (R2) in Theorem 4, §16.
- *Quantisation structure.* The theorem gives classical Maxwell-form transport. Quantisation of the gauge field (QED, photons as quantum excitations, Fock-space structure) requires separate quantisation procedure — the Hamiltonian admissibility derivation gives the Hamiltonian generator structure but the quantisation of the persistent gauge sector to full QED is not derived here.
- *Coupling constants.* The theorem gives the *form* of the dynamics (Maxwell U(1)) but not specific values of coupling constants (the fine-structure constant α, vacuum permittivity ε_0, etc.). These require separate derivation, addressed in other papers of the broader VERSF programme.
- *Phenomenological embedding.* The theorem gives the abstract gauge-sector dynamics; embedding in observable phenomenology (specific particle physics, classical electromagnetic phenomena, etc.) requires the matter coupling and additional structural input.
- *Logical derivation of (B1)–(B4) from BCB/TPB.* As detailed in §12.2, the four admissibility-class conditions are substrate-motivated by BCB/TPB but not strictly derived from them; the Main Synthesis Theorem's force inherits this status.

The Main Synthesis Theorem is therefore a *structural* result identifying Maxwell-form U(1) as the unique admissible continuum-limit reversible dynamics on the persistent sector under the stated conditions, not a derivation of conventional electrodynamics as a complete physical theory. The latter requires the open items above to be addressed in subsequent work.

---

15. Lorentz-Compatible Persistent Transport

The Main Synthesis Theorem (§14) establishes Maxwell-form U(1) gauge transport as the unique admissible reversible continuum dynamics on the persistent sector. But "Maxwell-form" carries an implicit content that deserves explicit treatment: conventional Maxwell theory is

Lorentz-covariant, with wave propagation at a universal speed c invariant across reference frames. Does the Main Synthesis Theorem's Maxwell-form result actually deliver Lorentz-covariant transport, or only the gauge-theoretic structural skeleton without the covariance content?

This section integrates two further companion-paper results — the *Wilson Limit paper* and the *Closure-Symmetry paper* — that bear directly on this question. The two papers together provide a plausible structural mechanism for why the Maxwell-form persistent transport is Lorentz-compatible at the substrate level, transforming the Main Synthesis Theorem from "Maxwell-form U(1) structure (covariance implicit)" into "Maxwell-form U(1) structure with substrate-level mechanisms supporting its Lorentz-compatibility."

15.1 The Lorentz-compatibility question

The Maxwell admissibility theorem (§12) gives Maxwell-form at $O(\varepsilon^0)$ under (B3) closure-geometry covariance, which inherits Lorentz covariance from the TPB causal bound in the continuum limit. But as §12.2 makes explicit, (B3) is *motivated by* TPB rather than strictly *derived from* it: the continuum Lorentz covariance is an additional structural input layered on top of TPB's finite-speed bound. The Main Synthesis Theorem inherits this status — its Lorentz-covariance content is conditional on (B3) holding, not derived from BCB/TPB alone.

A natural further question: are there substrate-level mechanisms that make (B3) *plausible* — or even, in suitable limits, *forced* — rather than being a separately-imposed admissibility postulate? If the substrate dynamics themselves favour Lorentz-compatible effective theory, then (B3) becomes a derived feature of the substrate rather than an external constraint. This is the question the Wilson Limit and Closure-Symmetry papers address.

15.2 The Wilson Limit paper: one-loop matching toward Lorentz compatibility

The *Wilson Limit paper* (cited: *The K = 7 Wilson Limit: One-Loop Matching, Lorentz-Compatible Effective Anisotropy, and the K = 7 Substrate Coupling*) addresses the following question. The bare substrate action on the K = 7 architecture admits, in principle, *anisotropic* couplings between spatial and temporal plaquette structures: β_s and β_t for spatial-plaquette and temporal-plaquette terms respectively. The dimensionless anisotropy parameter is:

$$\zeta := \sqrt{\beta_s / \beta_t}.$$

Lorentz-compatibility in the IR continuum limit requires $\zeta_{eff} = 1$ (or equivalently $\beta_s = \beta_t$ at the effective level): if $\zeta_{eff} \neq 1$, the effective theory has a preferred frame, with different propagation speeds in different directions, violating Lorentz covariance.

The Wilson Limit paper establishes a one-loop matching result for the flow of ζ under substrate refinement:

$$\zeta_{eff} - 1 = (1 - \delta)(\zeta_{bare} - 1) + \mathcal{O}(\text{higher-order corrections}),$$

with the one-loop suppression coefficient $\delta \sim 10^{-4}$, computed explicitly from a Hasenfratz–Karsch-style matching calculation on the K = 7 substrate. The interpretation: substrate refinement drives ζ toward the Lorentz-compatible fixed point $\zeta^* = 1$, with the bare anisotropy $\zeta_{bare} - 1$ suppressed by the factor $(1 - \delta) \approx 0.9999$ at the effective level. The fixed point $\zeta^* = 1$ corresponds to $c_{eff} = c_c$ (TPB bound, identified with the substrate's structural speed via Paper IV §4 of the broader programme), so Lorentz-compatible effective propagation is what the substrate dynamics drive toward.

Status and honest caveats. The Wilson Limit result is a *one-loop matching theorem*, not a full RG-flow theorem. The matching is single-scale (not an iterated recursion), which means the $(1 - \delta)$ factor is a one-step suppression, not an exponentially-suppressed convergence over many scales. The Wilson Limit paper itself is explicit about this:

- $\delta \sim 10^{-4}$ is too small to be the *primary* origin of empirical Lorentz precision (which experimentally is bounded at much tighter levels).
- Iterated RG-style convergence toward $\zeta^* = 1$ would require either multi-loop calculations beyond the one-loop matching, or substrate-level mechanisms that contribute additional suppression beyond the one-loop matching.
- The one-loop matching alone is therefore a *robustness theorem* — it shows that substrate refinement does not amplify bare anisotropy — rather than the *primary explanation* of empirical Lorentz-compatibility.

This honest framing matters: the Wilson Limit paper alone does not "derive" Lorentz invariance. It establishes that the substrate's matching dynamics are compatible with — and weakly favour — Lorentz-compatible effective theory.

15.3 The Closure-Symmetry paper: bare isotropy from spatial-temporal exchange

The *Closure-Symmetry paper* addresses the complementary question: does the bare substrate action already satisfy $\beta_s = \beta_t$ for structural reasons, prior to any matching dynamics? If so, the IR $\zeta^* = 1$ is reached *not* through suppression of a non-zero bare anisotropy (which would rely on the small $\delta \sim 10^{-4}$ Wilson-limit factor), but because $\zeta_{bare} = 1$ already at the substrate level.

The paper introduces a spatial-temporal exchange involution σ on the K = 7 architecture:

$$\sigma : \square_s \leftrightarrow \square_t,$$

mapping spatial-plaquette structures to temporal-plaquette structures and vice versa. The key structural claim is that σ is a *class-preserving symmetry* of the K = 7 hexagonal interface at the cardinality-matching level: under σ, the spatial-plaquette class and the temporal-plaquette class are interchanged in a structurally exact way, with equal cardinalities and matching combinatorial features.

If $\mathcal{S}_W$ (the Wilson sum of the substrate action) is invariant under σ, then the spatial-plaquette and temporal-plaquette terms must enter $\mathcal{S}_W$ with equal coefficients: $\beta_s = \beta_t$. The Closure-Symmetry paper's central result establishes σ-invariance of $\mathcal{S}_W$ at the cardinality-matching level, forcing $\beta_s = \beta_t = \beta_{K=7}$ exactly at the bare substrate action level (Theorem 2 of the Closure-Symmetry paper).

Combined with the Wilson Limit paper's matching result:

- From the Closure-Symmetry paper: $\zeta_{bare} = 1$ (exactly, at cardinality-matching level).
- From the Wilson Limit paper: $\zeta_{eff} - 1 = (1 - \delta)(\zeta_{bare} - 1) = (1 - \delta) \cdot 0 = 0$.

So $\zeta_{eff} = 1$ exactly at leading order: the IR fixed point is reached not by suppression but by bare isotropy, with the Wilson Limit's robustness theorem ensuring that the bare isotropy survives matching without being broken by one-loop corrections.

Status and honest caveats. The Closure-Symmetry paper is explicit about three levels of structural commitment for the σ symmetry:

- **Cardinality-level symmetry (proven):** σ is a bijection between the spatial-plaquette and temporal-plaquette classes at the cardinality level — they have equal cardinality and matching gross features.
- **Feature-matching (intermediate, partially established):** σ matches further structural features (orientation, incidence, multiplicity) beyond pure cardinality, strengthening the structural commitment toward a substantive symmetry.
- **Full chain-complex equivariance (open):** σ as an automorphism of the full substrate cell complex, with explicit equivariance under all relevant chain-complex operations. This level of structural commitment would establish σ as a substrate-physical symmetry rather than a feature-level matching.

The cardinality-level symmetry is sufficient for the substantive consequence ($\beta_s = \beta_t$ at the cardinality-matching level); the chain-complex realisation would strengthen this to full rigour. The Closure-Symmetry paper identifies the chain-complex realisation as the principal open structural item, with the cardinality-matching version being sufficient for the present-paper integration but the chain-complex realisation being the next combinatorial step required to fully close the Lorentz-compatibility story.

15.4 The combined picture: bare isotropy plus IR robustness

The Closure-Symmetry paper and the Wilson Limit paper together provide a coherent structural picture for the substrate-level origin of Lorentz-compatible Maxwell-form transport:

Layer	Mechanism	Result
Bare substrate	σ-invariance of $\mathcal{S}_W$	$\beta_s = \beta_t = \beta_{K=7}$ ($\zeta_{bare} = 1$)
Substrate refinement	One-loop matching (Wilson Limit)	$\zeta_{eff} - 1 = (1 - \delta)(\zeta_{bare} - 1)$, $\delta \sim 10^{-4}$
IR continuum	Composition	$\zeta_{eff} = 1$ exactly at leading order
Maxwell synthesis	Strong Synthesis + Main Synthesis	Lorentz-compatible Maxwell-form U(1) gauge transport

The architectural separation matters. The Closure-Symmetry paper provides the *bare isotropy* (a structural feature of the K = 7 architecture itself); the Wilson Limit paper provides the *IR robustness* (refinement dynamics do not amplify residual bare anisotropy). Together they account for Lorentz-compatibility at both the substrate-action level (bare) and the effective-theory level (IR), giving the (B3) condition of the Maxwell admissibility theorem a substrate-level grounding rather than leaving it as an externally-imposed admissibility postulate.

The crucial advantage of the combined picture: neither paper alone suffices.

- Wilson Limit alone gives only one-loop suppression with $\delta \sim 10^{-4}$, which is too weak to explain empirical Lorentz precision if bare anisotropy is order-unity.
- Closure-Symmetry alone gives exact bare isotropy at the cardinality-matching level, but does not address whether matching dynamics might break this isotropy at higher loops or under further substrate corrections.

Together, they provide the right structural architecture: bare isotropy from σ-symmetry, robust under one-loop matching, with the residual gap (whether higher-loop corrections or alternative substrate effects could break the isotropy at deeper levels) being the substantive open item.

15.5 The K = 7 unifying scale

A striking structural feature: the same K = 7 closure structure that controls the persistent transport sector (via the σ-family programme's master-action variation) and that provides the substrate-level Lorentz-compatibility (via the σ-symmetry of the Closure-Symmetry paper) also fixes the bare electromagnetic coupling.

From the *K = 7 paper* (cited: *The Fine Structure Constant from K = 7 Closure Geometry*), the bare coupling at the isotropic point $\beta_s = \beta_t = \beta_{K=7}$ is:

$$\beta_{K=7} = 2^7 \cdot (15/14) \approx 137.143.$$

This is structurally close to the empirical fine-structure constant inverse $\alpha^{-1} \approx 137.036$, with the small discrepancy attributed to higher-order corrections (the Wilson-limit one-loop matching contributing a suppression of order $\delta \sim 10^{-4}$, plus higher-loop and matter-coupling contributions). The σ-symmetry of the Closure-Symmetry paper fixes the *equality* $\beta_s = \beta_t$; the K = 7 closure-geometry derivation fixes the *value* $\beta_{K=7} = 2^7 \cdot (15/14)$.

Three structural roles for one number:

- $\beta_{K=7}$ is the bare electromagnetic coupling at the substrate level (K = 7 paper).
- $\beta_{K=7}$ is the σ-invariant common value of the spatial and temporal Wilson couplings (Closure-Symmetry paper).
- $\beta_{K=7}$ sets the scale of the one-loop matching suppression (Wilson Limit paper).

The same K = 7 closure structure controls all three roles. This is a tighter architectural connection than coincidence: the structural features of the K = 7 architecture (closure counting at

K = 7, the 14 = 2K loop-count decomposition, the closure-geometric β derivation) are the underlying source of all three.

15.6 Honest caveats and the chain-complex realisation gap

This is a synthesis section that integrates substantive results from companion papers; the caveats inherited from those papers carry forward into the present integration. The two principal caveats:

The chain-complex realisation of σ (load-bearing for the bare-isotropy mechanism). The σ symmetry of the Closure-Symmetry paper is established at the cardinality-matching level, with feature-matching partially established and full chain-complex equivariance flagged as open. The integration with the Wilson Limit paper relies on σ-invariance of $\mathcal{S}$_W at the cardinality-matching level, which is sufficient for the present synthesis but is the principal open structural item.

If the chain-complex realisation of σ can be explicitly constructed on a labelled K = 7 cell, the bare-isotropy mechanism is fully rigorous. If it cannot — i.e., if σ exists at the cardinality-matching level but does not lift to a full chain-complex automorphism — the bare-isotropy mechanism is still substantive but its strength is bounded by the cardinality-matching level. The chain-complex realisation is therefore the load-bearing open question for the substrate-level Lorentz-compatibility story.

The one-loop matching's limitations (load-bearing for the IR-robustness mechanism). The Wilson Limit result is single-scale matching at one loop. Multi-loop extensions, RG-flow analysis across multiple substrate-refinement scales, and explicit handling of substrate-corrections at higher orders are not addressed. The $(1 - \delta)$ factor with $\delta \sim 10^{-4}$ is a one-step robustness factor, not an exponentially-suppressed convergence. Empirical Lorentz-compatibility at much tighter precision than $\delta \sim 10^{-4}$ requires either (a) bare isotropy from the Closure-Symmetry mechanism (so that there is no anisotropy to suppress in the first place), or (b) multi-loop / multi-scale refinement of the Wilson Limit calculation, or (c) further substrate-level mechanisms not yet identified.

The combined picture handles (a) substantively at the cardinality-matching level. (b) and (c) remain open.

Summary of the integration. The Main Synthesis Theorem's Maxwell-form U(1) gauge transport result is now flanked by substrate-level mechanisms supporting its Lorentz-compatibility:

- Bare isotropy from σ-symmetry (Closure-Symmetry paper, cardinality-matching level).
- IR robustness from one-loop matching (Wilson Limit paper).
- Common K = 7 closure structure unifying the bare coupling value, the σ-symmetry, and the matching suppression scale.

The persistent gauge sector therefore carries not only Maxwell-form U(1) structure (Main Synthesis Theorem) but Lorentz-compatible Maxwell-form structure under the combined substrate-level mechanisms — modulo the chain-complex realisation gap and the one-loop matching limitations, which are flagged as the load-bearing open items.

This is a substantial strengthening of the Main Synthesis Theorem. It moves the framework from "Maxwell-form structure under (B3) Lorentz covariance as an admissibility postulate" to "Maxwell-form structure with substrate-level mechanisms supporting Lorentz-compatibility," with the (B3) admissibility condition itself now having substantive substrate-level motivation rather than being a separately-imposed external constraint. The (B3) condition is not "derived" in the strict sense (the chain-complex realisation gap and the one-loop matching limitations bound this), but its substrate-level grounding is substantially deeper than the (B1)–(B4) status of §12.2 alone would suggest.

PART IV — UNIFIED THREE-CLASS ADMISSIBILITY ARCHITECTURE

16. Theorem 4: Three Admissibility Classes and Their Coupling

Theorem 4 (Three admissibility classes generate three variational structures). *The σ-sector, the persistent gauge sector, and the record-formation sector of the broader VERSF programme belong to three structurally distinct admissibility classes — dissipative, reversible, and irreversible — each generating its own appropriate variational structure:*

- *Dissipative class → gradient-flow variational principle → σ-sector dynamics.*
- *Reversible class → unitary Hamiltonian generator (Wigner + Stone) → persistent gauge sector dynamics.*
- *Irreversible class → GKSL Lindbladian variational structure → record-formation sector dynamics.*

The unification programme for VERSF dynamics proceeds via inter-class coupling within this three-class architecture, not via uniform variational treatment.

Three routes for the inter-class coupling are identified. The enumeration is illustrative rather than exhaustive: (R1)–(R3) cover the three natural structural categories of inter-class coupling we have identified (intrinsic via shared configuration space; via a third sector that couples to both; via direct cross-term in the action), but further routes may be identifiable in subsequent work:

(R1) ***Intrinsic decoupling via the master configuration space.*** *The Hodge decomposition $C^1(C_6)$ = $Im(\delta) \oplus Ker(\delta)$* provides a natural composite configuration space in which the dissipative dynamics of the σ-sector (on Im(δ)) and the reversible dynamics of the persistent gauge sector (on Ker(δ*)) coexist as decoupled subsystems within the K = 7 architecture.*

(R2) ***Matter coupling for observable propagation.*** *The persistent gauge sector carries Hamiltonian Maxwell-form structure intrinsically (Main Synthesis Theorem) but acquires observable wave propagation, effective inertia, and matter-coupled phenomenology only through coupling to matter sectors elsewhere in VERSF. The structural analog is the conventional physics result that photons in vacuum carry the U(1) gauge structure intrinsically but acquire effective mass via matter coupling (e.g., the Meissner effect in superconductors). Care is needed in invoking this analog: in conventional QED, photons in vacuum already have propagating wave dynamics from the F^μν F_μν kinetic content; matter coupling adds effective mass on top of vacuum propagation. In the present framework, the K = 7 framework is 1D and its persistent sector Ker(δ*) has only one harmonic mode, so the "vacuum propagation" content of the photon analog has no K = 7 counterpart — matter coupling would have to supply both effective inertia and the spatial structure for wave propagation, not just the former. The higher-dimensional generalisations needed to make the photon analog substantive (with H^1 of higher rank carrying multiple modes and spatial-derivative kinetic terms) are an open structural problem (P3 in §22).*

(R3) ***Inter-class coupling within the K = 7 architecture.*** *Coupling terms in an extended master action that couple the dissipative, reversible, and irreversible sectors mediate energy and information exchange between the three admissibility classes. The specific form of such couplings within the K = 7 architecture is the subject of subsequent work.*

The discussion of the three routes mirrors the discussion in the sector-classification version of this paper, now enriched by the Main Synthesis Theorem of §14. Route (R2) is the structural analog of how photons in matter acquire effective mass via matter coupling — and is the natural physical content beyond the Maxwell-form admissibility result, providing the observable wave propagation and matter-coupled phenomenology that the persistent sector alone does not specify.

---

17. The Irreversible Record-Formation Class: GKSL Structure

The two admissibility classes identified so far (dissipative and reversible) do not exhaust the dynamical structures present in the broader VERSF programme. A third admissibility class — the *irreversible record-formation class* — corresponds to the commitment / decoherence dynamics that underlies tick events, measurement, and the irreversible recording of distinguishability into substrate records.

16.1 Characteristic features

The characteristic features of irreversible record-formation admissibility:

- **Non-unitary.** Unlike reversible admissibility, the dynamics does not preserve inner products; distinguishability content can be transferred from the system to the substrate's record sector, with the system's pure-state structure decohering into mixed states.
- **CPTP (Completely Positive Trace-Preserving) maps.** The dynamics is described by quantum operations that take density operators to density operators (rather than state

vectors to state vectors), reflecting the loss of pure-state structure under irreversible commitment.
- **Composable and continuous.** The dynamics is parametrised by a continuous time parameter, with composability and strong continuity (analogous to axioms (A1) and (A2) of the reversible class).
- **Entropy-increasing.** Unlike both the dissipative class (which decreases an admissibility-violation Lyapunov functional) and the reversible class (which conserves the Hamiltonian), the irreversible class increases von Neumann entropy as distinguishability content is transferred from system to record sector.

16.2 The GKSL theorem

By Lindblad's theorem (also independently due to Gorini, Kossakowski, and Sudarshan), any strongly continuous one-parameter semigroup of CPTP maps admits a generator in GKSL (Gorini–Kossakowski–Sudarshan–Lindblad) form:

$$d\rho/d\tau = -(i/\hbar)[H, \rho] + \Sigma_k (L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\}),$$

where:

- H is a self-adjoint operator (the Hamiltonian — same self-adjoint structure as the reversible class),
- $\{L_k\}$ are Lindblad operators encoding irreversible channels,
- $[\cdot,\cdot]$ is the commutator and $\{\cdot,\cdot\}$ is the anticommutator.

The GKSL generator decomposes into:

- **Antisymmetric (reversible) component:** $-(i/\hbar)[H, \rho]$, generating unitary rotation — the Hamiltonian structure of the reversible class.
- **Symmetric (irreversible) component:** the Lindblad dissipator $\Sigma_k (L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\})$, generating entropy increase via the Lindblad operators.

16.3 Connection to VERSF record formation

In VERSF terms, the Lindblad operators $\{L_k\}$ represent irreversible channels through which distinguishability content becomes physically committed to the substrate's record sector, producing persistent records and the asymmetry between past and future. This is the mechanism underlying tick events — the irreversible transitions that commit potentiality into actuality.

Measurement is not ontologically special in this framing. It is one instance of the general irreversible commitment structure. Whenever information becomes irreversibly recorded — in a laboratory detector, on a photographic plate, in the thermal jiggling of environmental molecules — the same GKSL mathematics applies.

The irreversible class therefore plays a distinctive role: it contains the *reversible class as a sub-structure* (the antisymmetric component) while extending it with the genuinely irreversible

Lindblad dissipator. The reversible class is the limit of the irreversible class as all $L_k \to 0$; the irreversible class is the natural completion of the reversible class to include record formation.

16.4 Status as a distinct admissibility class

The irreversible record-formation class is distinct from both the dissipative class and the reversible class:

- Distinct from the dissipative class: gradient flow on a Lyapunov functional is intrinsically deterministic and operates at the level of configurations ϕ, not density operators ρ. The dissipative class describes admissibility restoration of substrate configurations; the irreversible class describes information transfer from system to record sector.
- Distinct from the reversible class: unitary evolution preserves inner products; the irreversible class explicitly does not, via the Lindblad dissipator. The reversible class is the special case $L_k = 0$ of the irreversible class.

The three admissibility classes therefore form a hierarchy: dissipative (substrate admissibility restoration), reversible (system-internal distinguishability-preserving transport), and irreversible (system-to-record information transfer with entropy increase). Each class has its own characteristic dynamics, its own appropriate variational structure, and its own physical role within the broader VERSF programme.

---

18. The Unified Three-Class Architecture

The three admissibility classes together — dissipative, reversible, irreversible — form a unified architecture for VERSF dynamics. We summarise the architecture and its physical content.

17.1 The architectural table

Admissibility class	Variational structure	Characteristic dynamics	VERSF sector	Physical role
Dissipative	Gradient flow $\partial_\tau \phi = -\,\partial S/\partial\phi$	First-order, dissipative or invariant, never oscillatory	σ-sector (Im(δ))	Admissibility restoration
Reversible	Unitary Hamiltonian $U = e^{(-iH\tilde{\tau})}$	Conservative, oscillatory, distinguishability-preserving	Persistent gauge sector (Ker(δ*)), Maxwell-form in continuum	Reversible transport
Irreversible	GKSL Lindbladian $d\rho/d\tau = -(i/\hbar)[H,\rho] + \mathcal{D}[\rho]$	Non-unitary, entropy-increasing, system-to-record transfer	Record-formation sector	Commitment / decoherence

Each row of the table corresponds to a distinct mathematical structure, a distinct physical role, and a distinct set of derivations in the broader VERSF programme.

17.2 Coupling between classes

The three classes are *not* dynamically decoupled — they couple in physically meaningful ways:

- **Dissipative-Reversible coupling.** Within the master configuration space $C^1(C_6)$, the σ-sector dissipative dynamics operates on Im(δ) and the reversible Hamiltonian dynamics operates on Ker(δ*); under the Hodge decomposition $C^1(C_6) = \mathrm{Im}(\delta) \oplus \mathrm{Ker}(\delta^*)$ (framework paper Theorem 3, re-affirmed in §16(R1)), the two are dynamically decoupled as subsystems on orthogonal subspaces. The dissipative class does *not* set a "boundary condition" for the reversible class in any temporal sequencing sense — there is no point at which σ-sector dissipation "finishes" and "hands off" to reversible dynamics. Both operate simultaneously on their respective subspaces, with admissibility on Im(δ) and reversible Hamiltonian dynamics on Ker(δ*) coexisting as orthogonal features of the same configuration. The "coupling" between the two classes, in the intrinsic (R1) sense, is exactly this subspace coexistence within the same master configuration space; richer coupling (energy exchange between the classes) requires explicit inter-class coupling terms (R3) not present in the framework paper's master action.
- **Reversible-Irreversible coupling.** The Hamiltonian generator H of the reversible class is the antisymmetric component of the GKSL generator of the irreversible class. The reversible class is the L_k → 0 limit of the irreversible class; irreversibility is the addition of Lindblad channels to the reversible structure.
- **Dissipative-Irreversible coupling.** The substrate-level decoherence dynamics couples the irreversible record-formation processes to the substrate's admissibility structure. Records form *in* the substrate; the substrate's admissibility restoration (dissipative class) provides the physical medium for record formation (irreversible class).

These couplings together give the broader VERSF programme its complete dynamical architecture. The σ-sector framework, in this three-class architecture, occupies the dissipative-class corner; the Hamiltonian admissibility derivation and Maxwell admissibility theorem occupy the reversible-class corner; the record-formation work occupies the irreversible-class corner.

17.3 The σ-sector / persistent gauge sector unification, complete

The σ-sector / persistent gauge sector unification, originally posed as "how to upgrade the persistent gauge sector to Hamiltonian dynamics within the K = 7 master-action," is now positively resolved by the three-class architecture:

- The σ-sector and the persistent gauge sector belong to *different admissibility classes* (dissipative and reversible respectively).

- The persistent gauge sector carries Hamiltonian *generator structure* (not necessarily non-trivial dynamics — the eigenvalue ω is undetermined by Theorem 3 alone, with ω = 0 valid), from the reversible admissibility class, not from K = 7 catalogue penalties.
- In the continuum limit, the persistent gauge sector carries Maxwell-form U(1) gauge transport (Main Synthesis Theorem) as the unique admissible reversible dynamics under BCB/TPB and (B1)–(B4).
- The two sectors coexist as decoupled subsystems within the master configuration space $C^1(C_6)$ (R1), with coupling to the broader programme via matter sectors (R2) and inter-class structure (R3).
- The third admissibility class (irreversible record formation) provides the natural framework for the broader VERSF programme's tick events, measurement, and commitment dynamics.

The unification is therefore not a single-paper result but an *architectural recognition*: the σ-sector framework, the Hamiltonian admissibility derivation, the Maxwell admissibility theorem, and the irreversible record-formation work each contribute to a coherent three-class admissibility architecture within which the major dynamical sectors covered by these strands of work have their place. (Other sectors of the broader VERSF programme — matter sectors, gravity sector, possibly further classes not yet identified — may require additional architectural elements; the three-class architecture covers the three specific dynamical classes identified here, not necessarily every possible sector.)

CLOSING SECTIONS

19. The Refined Inertia Distinction: Hamiltonian Structure vs Intrinsic Mass

The earlier framing of the σ-sector / persistent gauge sector unification used "the inertia problem" to describe the absence of admissible mass-like terms from the K = 7 catalogue and the resulting lack of Hamiltonian oscillation on the persistent sector. Under the three-class architecture established in this paper, this phrasing requires careful refinement.

The correct distinction is between two separable structural concepts:

- **Hamiltonian structure** = the *reversible evolution generator* structure $U(\tilde{\tau}) = e^{-iH\tilde{\tau}}$ with self-adjoint H. This is the variational structure of the reversible admissibility class, derived via Wigner + Stone from the reversibility axioms (A1)–(A3).
- **Intrinsic mass / inertia** = the *restoring-force scale* that determines characteristic frequencies and provides observable propagation content. In Maxwell theory, the propagation content comes from the kinetic term ($F^{\mu\nu} F_{\mu\nu}$) and the matter coupling; intrinsic mass $m^2 A_\mu A^\mu$ would break gauge symmetry and is not present in the pure gauge sector.

These are different structural concepts. Hamiltonian structure can exist without intrinsic mass (the persistent sector under the Main Synthesis Theorem has Hamiltonian Maxwell-form structure without intrinsic mass). Intrinsic mass can appear in non-gauge sectors (matter sectors) where it doesn't break gauge symmetry.

Under the three-class architecture:

- **Hamiltonian structure on the persistent gauge sector** is established by Theorem 3 + the Hamiltonian admissibility derivation. The persistent sector has Hamiltonian structure intrinsically via the reversible admissibility class.
- **Maxwell-form U(1) transport on the persistent gauge sector** is established by the Main Synthesis Theorem. The persistent sector has *Maxwell-form* Hamiltonian structure in the continuum limit, not arbitrary Hamiltonian structure.
- **Intrinsic mass and observable propagation phenomenology** are *not* established by the framework alone; they come from coupling to matter sectors (R2 of Theorem 4). The Maxwell-form on the persistent sector is the *pure-gauge* dynamics; matter coupling provides effective mass and observable propagation.

The previous "inertia problem" framing conflated these. The refined statement: *the absence of K = 7-catalogue-admissible mass-like terms does not exclude Hamiltonian structure itself; it is consistent with the gauge sector's intrinsic masslessness in pure form, and effective mass / observable propagation arise from matter coupling — analogously to conventional physics, with the caveat noted in §16(R2)*. The persistent gauge sector of the σ-sector framework carries Maxwell-form U(1) Hamiltonian structure intrinsically (from reversibility + Maxwell admissibility); it would acquire effective mass and observable propagation through coupling to matter, analogously to how photons in vacuum carry U(1) gauge structure intrinsically and acquire effective mass through matter coupling (e.g., the Meissner effect). The qualifier matters: in conventional QED, photons in vacuum already propagate as waves (from the $F^{\mu\nu} F_{\mu\nu}$ kinetic content); in the present 1D K = 7 framework, the analog wave-propagation content requires higher-dimensional generalisations (P3) that lie beyond the present paper.

This refined distinction is much closer to the conventional physics picture than the previous framing suggested.

---

20. Implications for the Unification Programme

The Main Synthesis Theorem and the three-class architecture have substantive positive implications for the broader VERSF programme.

Implication 1: Positive resolution of the framework paper's open problem. The framework paper identified the upgrade to Hamiltonian dynamics on Ker(δ*) as the load-bearing open problem (P1). The Main Synthesis Theorem provides a substantively positive resolution: the persistent gauge sector carries Maxwell-form U(1) gauge transport as the unique admissible reversible continuum dynamics under BCB/TPB and (B1)–(B4). This is much stronger than the earlier negative result framing — it identifies the specific positive dynamics, not just rules out a class of approaches.

Implication 2: The three-class architecture unifies major VERSF strands. Three previously parallel strands of work — the σ-sector master-action framework, the Hamiltonian admissibility derivation, the Maxwell admissibility theorem (with its Strong Synthesis extension) — are now integrated into one architecture via the sector-classification recognition that they each describe a structurally distinct admissibility class. This is a substantial architectural unification at the level of the programme.

Implication 3: The σ-sector / persistent gauge sector unification has a substantive positive result at the framework level. Within the K = 7 framework, the σ-sector and the persistent gauge sector are now recognised as belonging to distinct admissibility classes, with the persistent sector carrying Maxwell-form structure in the continuum limit. The framework-level unification has a substantive positive resolution — *Maxwell-form U(1) structure on the persistent sector under the stated conditions* — with three things to keep in scope:

- The substantive Maxwell content awaits higher-dimensional generalisations to be fully realised (per the discussion in §14 of what "Maxwell-form" means on the 1D K = 7 specialisation; the substantive wave-propagation content lives in the higher-dimensional version).
- The foundational dependencies — the Strong Synthesis Theorem, the BCB/TPB-to-(B1)–(B4) admissibility-class motivation (per §12.2), and the K = 7-to-general-substrate embedding — require the broader programme's results to be in place.
- The implementation work — matter coupling for observable propagation, quantisation to full QED, coupling constants, phenomenological embedding — is substantive subsequent work, well-posed within the architecture established here but not contained within this paper.

The unification is therefore "positively resolved" rather than "positively closed": substantial framework-level progress, with implementation and foundational dependencies still in scope rather than concluded.

Implication 4: The Maxwell-form result is conditional on BCB/TPB and the admissibility class (B1)–(B4). The Main Synthesis Theorem inherits its conditions from the Maxwell admissibility theorem. If BCB or TPB fail in some regime, or if the admissibility class (B1)–(B4) needs modification, the Main Synthesis Theorem fails in that regime. The substrate axioms are therefore the load-bearing structural inputs, and their derivation in the broader programme is what underwrites the Main Synthesis Theorem.

Implication 5: Matter coupling is the next major step. Under the three-class architecture, the next natural step is the matter-coupling derivation: deriving how matter sectors couple to the persistent gauge sector, providing effective inertia and observable propagation phenomenology. This is well-posed within the architecture and is the structural analog of how conventional electromagnetism couples to charged matter.

Implication 6: The third class (irreversible record formation) is naturally positioned for further development. With the dissipative and reversible classes now integrated via the present paper, the third class (irreversible record formation) is the natural site for further architectural development — specifically, how record formation couples to the dissipative and reversible classes, and what specific Lindblad structure emerges from VERSF's commitment dynamics.

---

21. What This Establishes — and What It Does Not

What this paper establishes:

1. The σ-sector and the persistent gauge sector of the σ-sector framework are most naturally assigned to two structurally distinct admissibility classes — dissipative and reversible (Lemma 1, Theorem 2, Theorem 3, with the assignments themselves treated as structural postulates per §3).
2. The dissipative admissibility class's variational structure (gradient flow on the $K = 7$ catalogue) correctly governs the σ-sector but does not generate mass-like terms (Lemma 1, Theorem 2).
3. The reversible admissibility class's variational structure (unitary Hamiltonian generator, derived via Wigner + Stone) governs the persistent gauge sector, preserving the Wilson-loop topological invariant (Theorem 3).
4. The persistent gauge sector $\mathrm{Ker}(\delta^*)$ of the $K = 7$ framework is the $K = 7$ specialisation of the broader substrate's refinement-stable observable sector $H^1(G(\Lambda))$ (§11).
5. Under BCB, TPB, and the admissibility class (B1)–(B4), the unique refinement-stable continuum-limit reversible dynamics on the persistent sector is Maxwell-form U(1) gauge transport at $O(\varepsilon^0)$ — the Main Synthesis Theorem (§14), established by composition of Theorem 3 with the Hamiltonian admissibility derivation and the Strong Synthesis Theorem.
6. The three admissibility classes (dissipative, reversible, irreversible) form a unified three-class architecture for VERSF dynamics (Theorem 4, §18), with the σ-sector framework occupying the dissipative-class corner and the persistent gauge sector occupying the reversible-class corner.

What this paper does not establish:

- *Full relativistic completion.* The Main Synthesis Theorem gives Maxwell-form U(1) at $O(\varepsilon^0)$; higher-order corrections, non-Abelian extensions, and full relativistic field theory beyond Maxwell are not addressed.
- *Matter coupling.* The pure-gauge-sector dynamics is established; coupling to charged matter (electrons, sources, currents from non-substrate-level matter sectors) requires separate derivation.
- *Quantisation structure.* Classical Maxwell-form transport is established; quantisation of the gauge field to full QED (photons as quantum excitations, Fock-space structure, second-quantisation procedures) is not derived.
- *Coupling constants.* The *form* of the dynamics is established; specific values of coupling constants (fine-structure constant, vacuum permittivity, etc.) are not derived here (though they are addressed in other papers of the broader VERSF programme).

- *Phenomenological embedding.* The abstract gauge-sector dynamics is established; embedding in observable phenomenology (specific particle physics, classical electromagnetic phenomena, observable predictions) requires matter coupling and additional structural input.
- *Higher-dimensional generalisations.* The K = 7 framework is 1D on the rim cycle. Higher-dimensional generalisations of the K = 7 architecture, where richer non-Abelian gauge structure becomes available, are not addressed.
- *Independent derivation of the companion papers' results.* The Main Synthesis Theorem composes three established results from companion papers; it does not re-derive them. If those companion papers' derivations are revised, the composition would need revision accordingly.
- *Independent derivation of BCB and TPB.* These substrate axioms are taken from the broader VERSF programme; their derivation is in those papers, not here.

The Main Synthesis Theorem is therefore a *structural composition* result. Its strength comes from the combination of three established results, plus the sector-classification recognition that makes the composition possible. Its limitations come from the limitations of its component pieces and the conditions under which they apply.

---

22. Open Problems

P1 (load-bearing for observational contact). Matter coupling derivation for the persistent gauge sector. The Main Synthesis Theorem establishes Maxwell-form U(1) pure-gauge dynamics; matter coupling provides observable propagation, effective inertia, and phenomenology. Which VERSF matter sector couples naturally? What is the form of the coupling? Does it reproduce conventional gauge-matter coupling phenomenology in appropriate limits?

P2. Quantisation of the persistent gauge sector to full QED. The Hamiltonian admissibility derivation establishes the Hamiltonian generator structure; second-quantisation procedures producing photon Fock space, particle creation/annihilation operators, and the full QED apparatus need to be developed within the VERSF setting.

P3. Higher-dimensional lifting of the K = 7 architecture. The K = 7 framework is 1D on the rim cycle; richer non-Abelian gauge structure (SU(2), SU(3), Yang-Mills theory) presumably requires higher-dimensional K_N architectures with richer cohomology. What are the natural generalisations?

P4. Specific form of inter-class coupling (R3) within the K = 7 architecture. What is the natural form of a coupling term that mixes the dissipative, reversible, and irreversible variational structures?

P5. Explicit substrate-to-K = 7 embedding. The Main Synthesis Theorem assumes the K = 7 framework is the K = 7 specialisation of the broader substrate framework on which the Maxwell admissibility theorem operates. The explicit embedding — showing that K = 7 satisfies BCB/TPB and (B1)–(B4) as a specific substrate — needs detailed development.

P6. Detailed derivation of the irreversible record-formation class's Lindblad operators within VERSF. The general GKSL form is summarised in §17; specific {L_k} for VERSF's commitment dynamics need to be derived.

P7. Composition of the three classes in fully coupled dynamics. The Main Synthesis Theorem gives the reversible class on the persistent sector; the full programme requires coupled three-class dynamics for the complete VERSF picture.

P8. Coupling constants and dimensional analysis. The Maxwell-form result gives the structure; specific coupling constants (fine-structure constant α, electromagnetic coupling, etc.) need separate derivation.

P9. Phenomenological matching. Demonstrating that the Main Synthesis Theorem's continuum-limit Maxwell-form U(1) matches conventional electromagnetism in appropriate limits — including specific phenomenological predictions that could be tested observationally.

P10. Extension to non-Abelian gauge sectors. SU(2), SU(3), and full Standard Model gauge structure require non-Abelian extensions of the persistent cohomological transport framework. The broader VERSF programme has work on this (deriving the Standard Model gauge group from BCB distinguishability geometry); integrating it with the three-class architecture is a substantive open task.

---

23. Conclusion

The σ-sector master-action framework paper established the kinematic structure of the master configuration space $C^1(C_6)$ on the rim cycle of the K = 7 architecture, with the orthogonal Hodge decomposition into the σ-sector subspace Im(δ) and the persistent gauge sector subspace Ker(δ*) ≅ $H^1(C_6)$. Under uniform gradient-flow variational treatment of both subspaces, the framework's persistent sector was static, and the upgrade to non-trivial dynamics on Ker(δ*) was identified as the load-bearing open problem.

This paper has resolved the open problem positively through a two-stage process. The first stage was the *sector-classification* recognition that the σ-sector and the persistent gauge sector are most naturally assigned to two structurally distinct admissibility classes — dissipative and reversible — each with its own appropriate variational structure. The framework paper's static persistent sector under uniform treatment was not a misapplication or error: the framework paper explicitly flagged the Hodge heat flow everywhere as one specific postulate among several, with the Hamiltonian alternative on Ker(δ*) noted as available. The present paper makes the alternative substantive by sourcing it from the reversible admissibility class — drawing on the Hamiltonian admissibility derivation work via Wigner + Stone — so that the persistent gauge sector carries Hamiltonian generator structure from the reversibility principles rather than from K = 7 catalogue penalties.

The second stage was the *composition with established results* in the broader VERSF programme. The Hamiltonian admissibility derivation establishes that reversible distinguishability-preserving evolution necessarily takes unitary form $U(\tilde{\tau}) = e^{(-iH_h \tilde{\tau})}$. The Maxwell admissibility uniqueness theorem establishes that under BCB and TPB, within the admissibility class (B1)–(B4), Maxwell-form U(1) gauge transport is the unique continuum-limit transport theory at $O(\varepsilon^0)$. The Strong Synthesis Theorem combines these with the substrate-level identification of $H^1(G(\Lambda))$ as the unique refinement-stable observable sector. Composing all three: under BCB/TPB and (B1)–(B4), the σ-sector framework's persistent gauge sector — considered as the K = 7 specialisation of $H^1(G(\Lambda))$ — carries Maxwell-form U(1) gauge transport as the unique admissible reversible continuum dynamics. This is the Main Synthesis Theorem of the paper.

Part IV completed the architecture by integrating a third admissibility class — the irreversible record-formation class governed by GKSL Lindbladian dynamics — corresponding to the commitment / decoherence sector of the broader VERSF programme. The three classes together (dissipative, reversible, irreversible) form a unified three-class admissibility architecture in which each major sector of the VERSF programme has its place.

The reframing has consequences beyond the immediate question. The continuum-limit paper's four structural obstacles included two — parabolic vs hyperbolic, dissipative vs conservative — that were originally posed as obstacles to be resolved by a Hamiltonian upgrade within K = 7. Under the three-class architecture, these are now *category distinctions* between structurally distinct admissibility classes rather than conflicts to be resolved within one. The σ-sector is parabolic and dissipative because it belongs to the dissipative class; the persistent gauge sector is hyperbolic and conservative (carrying Maxwell-form transport) because it belongs to the reversible class.

The σ-sector / persistent gauge sector unification programme — which the framework paper flagged as the load-bearing structural choice between postulates governing dynamics on $Ker(\delta^*)$ — is therefore substantively positively resolved at the framework level. The persistent gauge sector carries Maxwell-form U(1) gauge transport as the unique admissible reversible continuum dynamics; the dissipative class correctly describes the σ-sector; the inter-class coupling structure (R1, R2, R3) accommodates the σ-sector / persistent gauge sector relationship via matter coupling, intrinsic decoupling, and inter-class coupling within the K = 7 architecture.

What this paper has therefore established is a *synthesis bridging three previously parallel strands of the VERSF programme*: the σ-sector master-action framework, the Hamiltonian admissibility derivation, and the Maxwell admissibility theorem (with its Strong Synthesis extension). The bridging composition gives a substantively positively resolved framework-level unification, with Maxwell-form U(1) gauge transport identified as the unique admissible continuum dynamics on the persistent sector under the stated conditions, and with implementation and foundational dependencies (per §20 Implication 3) still in scope rather than concluded. The remaining open problems — matter coupling for observable propagation, quantisation to full QED, coupling constants, phenomenological embedding, higher-dimensional generalisations, and the detailed integration of the third admissibility class — are now well-posed within the three-class architecture, ready for substantive subsequent work.

The substantive conceptual content of the paper: the σ-sector, persistent gauge sector, and record-formation sector are three structurally distinct admissibility classes, each with its appropriate variational structure, together forming a unified architecture in which Maxwell-form U(1) gauge transport on the persistent sector emerges as the unique admissible reversible continuum dynamics under the substrate axioms of the broader VERSF programme. This is the positive resolution of the framework paper's open problem, and the architectural unification of three major strands of the broader VERSF programme.