The Structure Group of the Maxwell Connection

One Circle, Not Two — Discharging the Gauge Identification

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The question

The charge paper in this programme (*Charge Quantization from Compact Phase*) argued that electric charge comes in whole-number steps because the phase of physics is a circle — a dial that comes back to its start after one full turn — and a gear meshed to a circular dial can only turn a whole number of times per revolution. That argument rested on exactly one named assumption: that the dial electromagnetism uses and the dial quantum interference uses are *one and the same dial* — one circle, not two. The charge paper called this the Gauge Identification, carried it openly as a condition, and named the present paper as the place where it would be checked.

What this paper does

This paper performs the check. The programme's Maxwell-sector papers build electromagnetism out of the substrate's transport structure — the same structure whose phase the U(1) paper proved to be a circle. So the check is not a comparison of two separately built dials to see whether they happen to match. It is more like opening the casing of a machine and following the linkage: if the electromagnetic dial is driven by the same spindle as the transport dial — if it *is* that spindle, viewed from the electromagnetic side of the machine — then there is nothing to identify, because there was never a second dial. The paper exhibits the linkage step by step, and then settles two further matters that a careless check would skip: that the dial really is a circle rather than an endlessly accumulating counter (the difference between a clock face and an odometer), and that the two descriptions agree not just in shape but point for point, so that "one full turn" means the same thing on both sides.

The honest condition

The check is a check *against* the Maxwell-sector papers, and it consumes three specific results from them, each stated in this paper in exact form with a marked slot for its citation. If those papers establish the three results as stated, the Gauge Identification is discharged: charge quantization then rests only on the programme's foundations, with no free assumption left in the chain. If any of the three holds only in a weaker form, this paper says precisely which step of the check fails, and the charge paper's conclusion reverts — undamaged — to its original conditional standing. Nothing is claimed beyond what the cited results support.

What follows if the check succeeds

The oldest discreteness in physics — the ladder of electric charge — then traces in an unbroken chain to the programme's starting commitments: a finite world that cannot sustain structure nothing in it could witness is forced to a circular phase; electromagnetism, built from transport, uses that circle because it is made of it; and a circle admits only whole-number gear ratios. The ladder of charge is the meshing of one gear on one circle.

Abstract

The charge paper derives the integer charge lattice from the compact transport phase of the U(1) paper, conditional on exactly one new premise: the Gauge Identification (GI), the claim that the structure group of the Maxwell connection is the derived compact $\mathbb{T}$ — that the electromagnetic gauge phase *is* substrate transport holonomy, not a second circle and not the line $\mathbb{R}$. The charge paper's §3.2 audit narrows GI to an exact disjunction and names the discharge obligation in concrete form: exhibit the Maxwell-sector phase as transport holonomy and read off its structure group. The present paper performs that reading.

The reading consumes three results from the programme's Maxwell-sector papers, each stated here in exact form with a marked citation slot and a named fallback if the source establishes a weaker form (§3): **M1 (Transport Origin)** — the Maxwell connection's parallel transport along an admissible path *is* the substrate transport along that path, identity by construction rather than correspondence between independently built structures; **M2 (Holonomy Exhaustion)** — the gauge-invariant content of the Maxwell sector is exhausted by closed-loop holonomies, with the field strength as their infinitesimal form; **M3 (No Independent Lift)** — no Maxwell-sector result requires the accumulated real phase θ, as opposed to the unit-modulus holonomy e^{iθ}, as admissible invariant content, and no construction treats total winding as physical structure of the holonomy itself at any step — real-valued flux through realized spanning surfaces being configuration-level curvature records inside M2's algebra, not lifts. The imports' asymmetry is calibrated (§3.1): M1 carries the discharge weight — it determines whether there are one or two phase structures at all, identity-by-construction replacing identification-by-correspondence — while M2 carries the sole GI-striking fallback and M3 the cross-arc consistency risk. Given M1–M3 and the inherited compact-phase result, three steps complete the discharge. (i) The **exhibition** (§4): by M1 the Maxwell holonomy assignment and the substrate transport holonomy assignment are one map, quantified over admissible paths without fixing the value any path carries — assignment-neutral, preserving the charge paper's §12 independence (§7). (ii) The **structure-group reading** (§5): the structure group is fixed by which holonomies count as the same holonomy; by the U(1) paper's exact periodicity the admissible object is e^{iθ} with θ and θ + 2π one point, a branch selection being an unwitnessed distinction; M3 confirms the Maxwell

sector adds no winding record; an unrealized winding record cannot ground a distinction under the realized-witness discipline; hence the structure group is $\mathbb{T}$, not $\mathbb{R}$ — the odometer reading that conventional electromagnetism leaves undetermined, here forced [Proven, conditional on M1–M3 and inherited inputs]. (iii) The **canonicality of the identification** (§6): the continuous automorphisms of $\mathbb{T}$ are exactly $z \mapsto z$ and $z \mapsto \bar{z}$, and the non-injective endomorphisms $z \mapsto z^k$ ($|k| \geq 2$) are excluded by M1's identity-by-construction; the residual freedom is orientation alone — the sign convention of charge, physically empty — so the 2π-identification consumed by the charge paper's §6 is the *same* identification, not an isomorphic copy, and the integer lattice transfers verbatim.

Consequences are tracked exactly (§8): GI's marker moves, upon verification of the three slots, from [Conditional] to [Proven, conditional on M1–M3 [Inherited] and the inherited compact-phase result]; the charge lattice sheds its one free premise, its conditionality relocating — not vanishing — into the Maxwell arc's source statuses; the Assignment-conditional strengthened refuter of the charge paper's §13 is absorbed; and the empirical off-lattice channel becomes a single falsification lever bearing jointly on three arcs of the corpus. The refuters of the present reading are named (§9): a non-transport component of the Maxwell connection (against M1), a realized winding witness (against M3 — which would additionally strike the U(1) paper upstream), and source-verification failure at any of the three slots, each with its stated fallback. The paper adds no new conditional to the corpus; it removes one, if and only if the three imports verify as stated.

Epistemic markers: [Inherited] imported from prior VERSF papers; [Imported-External] imported from standard mathematics outside the programme, carried at the external source's standing; [Proven] established here; [Conditional] holding under stated inputs; [Conjectural] motivated but unproven; [Open] undecided.

Table of Contents

---

1. Introduction — The Obligation and Its Shape

The charge paper closes with one conditional standing between its result and the programme's foundations: the Gauge Identification, GI — that the phase to which electric charge couples is the substrate transport phase derived in the U(1) paper, one compact circle and not two. The charge paper does not merely carry the premise; it specifies its own successor. The §3.2 audit there converts GI from a plausibility claim into an exact disjunction — either electromagnetism uses the derived transport circle, or it requires a second physically independent circle whose admissible invariant content must be exhibited — and states the discharge obligation in a single sentence: *exhibit the Maxwell-sector phase as transport holonomy and read off its structure group — nothing more, nothing vaguer.* The present paper is that exhibition and that reading.

The shape of the task is worth fixing before the work begins, because it determines what kind of paper this is. It is not a derivation paper: no new physics is derived here, and the corpus's stock of conditionals is meant to shrink by one, not grow. It is an *audit* — a verification of consistency between two arcs of the corpus that were built for different purposes: the phase arc, which derives the compact U(1) from finite distinguishability, and the Maxwell arc, which constructs the electromagnetic sector from admissibility constraints on substrate transport. GI is the statement that the U(1) appearing in the second arc is the U(1) derived in the first. An audit succeeds or fails against its sources, and this one names its sources exactly: three results it consumes from the Maxwell-sector papers, stated in §3 in the precise form the argument requires, each with a marked citation slot and a named fallback position if the source establishes something weaker. The paper's own logic — that M1–M3 plus the inherited compact-phase result discharge GI — is established here unconditionally; whether M1–M3 hold as stated is the audit proper, and the paper is complete exactly when its three slots are filled or its fallbacks invoked.

Two subtleties give the audit its content, and naming them up front prevents the appearance that the reading is trivial. The first is the **odometer question** (§5). Conventional electromagnetism, written with a real-valued potential, leaves undetermined whether the gauge group is the circle $\mathbb{T}$ or the line $\mathbb{R}$ — whether phase is a clock face, on which twelve and twenty-four hours are one position, or an odometer, which records total accumulation and never returns. Observables involving quantized charge cannot distinguish the two, which is exactly the circularity the charge paper diagnoses: compactness implies quantization, quantization is the evidence for compactness, and the explanation chases its tail. The reading must therefore do more than exhibit a connection; it must establish that the connection's structure group is the circle and not the line, and it must do so from the substrate side, where the U(1) paper's machinery has independent force. The second is the **canonicality question** (§6). Two circles are always abstractly isomorphic; an audit that ended at isomorphism would have established nothing, since every circle matches every other in shape. The discharge requires that the identification be canonical — fixed by the construction itself, point for point, so that the 2π-periodicity the charge paper's §6 converts into integrality is the *same* identification on both sides and not a matching copy with a hidden winding between them. Both subtleties resolve, and the resolutions are the paper's technical content; but they resolve because of specific features of the Maxwell-sector construction, which is why the imports of §3 must be stated exactly rather than gestured at.

2. The Premise to Discharge — GI in Structure-Group Form

The premise is imported in the form the charge paper fixes [Inherited, charge paper §3]:

Gauge Identification (GI). The gauge phase of electromagnetism is the substrate transport phase: the phase to which charge couples ranges over the same compact circle derived in the U(1) paper, with the same exact 2π-periodicity. Equivalently, in gauge-theoretic terms: the structure group of the Maxwell connection is the derived compact $\mathbb{T} = \{ z \in \mathbb{C} : |z| = 1 \}$, and the connection's holonomy is substrate transport holonomy — not a second circle, and not the line $\mathbb{R}$.

Three components of the claim require separate verification, and the paper's central sections track them one each:

(GI-a) Identity of holonomy. The Maxwell connection's parallel transport is substrate transport — the holonomy the electromagnetic sector assigns to an admissible path is the transport holonomy of that path. Verified in §4 from M1.

(GI-b) Compactness of the group. The group in which that holonomy is valued — the structure group — is the compact circle $\mathbb{T}$, not its universal cover $\mathbb{R}$. Verified in §5 from M2, M3, the inherited exact-periodicity result, and the realized-witness discipline, with the §4 identity (M1) effecting the transfer.

(GI-c) Canonicality of the identification. The identification between the electromagnetic circle and the transport circle is the identity map up to orientation — not merely an isomorphism, and not a k-fold winding. Verified in §6 from M1 and the classification of continuous endomorphisms of $\mathbb{T}$.

The conjunction (GI-a) ∧ (GI-b) ∧ (GI-c) is GI in full strength: the strength the charge paper consumes. A discharge of (GI-a) and (GI-b) without (GI-c) would leave the charge lattice intact as a lattice but would leave open a relabeling of its rungs; the charge paper's §10 disclaims the unit in any case, but a non-canonical identification of $|k|$-fold type would do worse than relabel — it would make the electromagnetic phase strictly coarser than transport phase, a genuine non-identity. §6 closes this.

What this paper does *not* re-litigate: the derivation of the compact transport phase itself, consumed whole from the U(1) paper at its source status [Inherited]; the representation-theoretic chain from compact phase to integer lattice, consumed whole from the charge paper [Inherited]; and the motivational case for GI — the four structural coincidences of the charge paper's §3.2 — which the discharge supersedes rather than repeats. Motivation is what a premise has before its audit; this paper is the audit.

3. The Maxwell-Sector Imports — Statements, Slots, and Fallbacks

The audit consumes three results from the Maxwell-sector papers. Each is stated here in the exact form the argument requires; each carries a citation slot, marked ⟨cite: …⟩, to be fixed against the source papers; and each carries a fallback — the position the corpus occupies if the source establishes the result only in a weaker form. The fallbacks are not rhetorical insurance. They are the audit's failure modes stated in advance, in the same discipline by which the charge paper stated its refuters: a verification whose failure conditions are unnamed is not a verification.

M1 (Transport Origin). *The Maxwell connection is constructed as substrate transport: its parallel transport along an admissible path is the substrate transport along that path, and the electromagnetic gauge potential is the local description of substrate transport phase accumulation. The relation is identity by construction — the Maxwell sector does not build an independent phase structure and then relate it to transport; it presents transport in electromagnetic form.* [Inherited — ⟨cite: Maxwell-sector paper, construction section⟩]

Verification obligation: the source must establish construction-as-transport, not merely construction-from-transport-constraints. The distinction is load-bearing. A sector built *from* admissibility constraints on transport might still introduce its own phase variable, constrained by transport but not identical to it; a sector built *as* transport has no second variable to identify. The verifier should locate the step in the source at which the electromagnetic phase is introduced and confirm that it is introduced as (a function presenting) the transport phase, not as a new degree of freedom.

Fallback if weaker: if the source establishes only construction-from-constraints, GI reverts to [Conditional], the gap relocates to the single named step — the introduction of the electromagnetic phase variable — and the discharge obligation sharpens to: prove that step's variable is transport phase. Everything in §§5–6 of the present paper survives as conditional machinery awaiting that proof.

M2 (Holonomy Exhaustion). *The gauge-invariant content of the Maxwell sector is exhausted by closed-loop holonomies: every admissible electromagnetic invariant is a function of the holonomies $h(\gamma)$ of closed loops γ — the transport around γ — with the field strength as their infinitesimal form, curvature as the local density of loop phase. Where a potential A is locally defined, the holonomy carries the familiar gloss $h(\gamma) = \exp(i\oint_\gamma A)$ with $F = dA$; the import is stated potential-free so that it does not quietly presuppose a globally defined A.* [Inherited — ⟨cite: Maxwell-sector paper, invariance/observables section⟩; the general principle that gauge-invariant content is closed-loop content is [Imported-External], carried as in the Assignment Paper]

Verification obligation: the source must not posit electromagnetic invariants outside the holonomy algebra — in particular, no invariant local phase and no invariant open-path phase.

Fallback if weaker: any invariant outside the holonomy algebra is precisely a candidate for the charge paper's §13 refuter — admissible electromagnetic content not representable as transport holonomy — and the audit converts from a discharge into a refutation procedure: exhibit the invariant, test it against transport holonomy, and report. This is the one fallback that could *strike* GI rather than merely suspending its discharge, and it is named as such.

M3 (No Independent Lift). *No Maxwell-sector result requires the accumulated real phase θ — as opposed to the unit-modulus holonomy e^{iθ} — as admissible invariant content. Further, no Maxwell-sector construction treats total winding — the integer separating θ from $\theta + 2\pi n$ — as physical structure of the holonomy itself at any step, including non-invariant intermediaries. Real-valued magnetic flux through a realized spanning surface is not an instance and is not forbidden: it is configuration-level curvature data carried by a realized record, inside M2's algebra (§5, step three), distinguishing field configurations, never fiber points. The real-valued potential and the bare line integral $\oint A$ — absent such a record — are descriptive bookkeeping, gauge-covariant intermediaries; the invariant object at the fiber is the exponential.* [Inherited — ⟨cite: discharged one-hop via the companion paper *From Real Circulation to Compact Holonomy*, which performs the construction-level check against the Maxwell-sector papers; the companion itself consumes M2, so the M2 slot is to be fixed first or jointly with this one⟩]

Verification obligation: the verifier should search specifically for any construction that compares holonomies across the 2π boundary as if they were distinct. This is the check of the statement's construction-level clause; the invariant-level clause is covered by M2's verification, per the scope paragraph below.

Fallback if weaker: a Maxwell-sector commitment to the lift would set the structure group to $\mathbb{R}$, contradicting not only GI but the U(1) paper's exact-periodicity result upstream — the corpus would carry an internal inconsistency between two arcs, to be resolved at source, with this paper suspended until it is. This fallback is an inconsistency alarm, not a graceful degradation, and is flagged accordingly.

The relation between M2 and M3 — scope, not redundancy. On their face the two imports nest: a function of the exponential e^{iθ} is automatically blind to θ against $\theta + 2\pi$, since the exponential has already quotiented by the winding — so if every invariant is a function of loop holonomies (M2), no invariant distinguishes lifts, which is M3's conclusion. The imports are nonetheless carried separately, because their quantifiers differ. M2 quantifies over *invariants*: it fixes what the sector's gauge-invariant content is. M3 quantifies over *results and constructions*: it forbids the sector from treating winding as physical structure anywhere in its machinery — including non-invariant intermediate steps that might smuggle a winding record into the catalogue by construction, before the question of invariance is ever posed. A sector could satisfy M2 at the invariant level while a construction inside it compared holonomies across the 2π boundary as if they were distinct, and that comparison — not any invariant — is what M3's verification obligation hunts. The entailment therefore runs only at the invariant level; at the construction level M3 is independent, and the inconsistency alarm of its fallback attaches to the construction-level clause. A verifier who confirms M2 has confirmed M3's invariant half for free; the construction-level half remains a separate check, and the import count is three because the checks are three.

One further import is consumed silently throughout and is recorded here once: the inherited compact-phase result of the U(1) paper, in the structural form the charge paper's §2 fixes — $\mathbb{T}$ compact, connected, one-dimensional, abelian; exact periodicity as admissible structure; the Admissibility Conservation Lemma; the realized-witness discipline of the closure argument [all Inherited, at source statuses].

3.1 Why M1 Carries Most of the Weight

The three imports do not contribute equally, and the asymmetry should be stated before the audit consumes them. M2 and M3 determine how the electromagnetic phase behaves once it exists — which invariants it carries, whether it lifts. M1 determines whether there are one or two phase structures in the first place, and that is a difference in kind rather than in degree.

The counterfactual makes the asymmetry exact. Suppose M2 and M3 held while M1 failed: the electromagnetic phase would be a compact holonomy structure carrying no independent lift — and would remain a distinct object from substrate transport. Every step of §§5–6 that runs on M2 and M3 would still run, and the central question of the charge paper would survive untouched: why should those two circles be identified? The discharge would be exactly where it started. M1 is what removes the question rather than answering it: if the Maxwell connection is substrate transport, identity-by-construction replaces identification-by-correspondence — the electromagnetic phase is not a second structure matching the transport phase but the transport phase viewed through electromagnetic variables, and GI ceases to be a bridge between two circles because there is only one circle. M2 and M3 fix the structure group's form; M1 fixes whether there are one or two structure groups available to discuss.

Why identity-by-construction is stronger than correspondence. The strength of M1's form deserves its own statement, because it is where the discharge's security comes from. A correspondence argument compares two independently constructed structures and shows that they agree — and such arguments are permanently vulnerable in a way no amount of checking fully closes: to hidden relabellings, to alternative parameterizations, to undiscovered intermediate structure sitting between the two objects compared. Identity-by-construction is a different kind of claim. It does not show that two structures agree; it shows that only one structure was ever constructed — there is nothing for a relabelling to sit between. This is the deep reason M1 carries the discharge weight: if the Maxwell phase is introduced as transport phase from the beginning, the identification questions a correspondence would have to answer one by one never arise. They are not solved; they dissolve. The paper's one concrete instance makes the general point exact: the k-fold winding scenario of §6 — the hidden relabelling par excellence, a power map $z \mapsto z^k$ lurking between two circles — is excluded not by an argument against each k but by the absence of a second circle for the map to target. Correspondence would have owed that argument; identity owes nothing — though what identity does owe, it owes in one place: the entire weight of the claim concentrates in the M1 slot, which is why the verifier's effort is allocated there first and hardest. The burden is not smaller; it is relocated — the same doctrine §8 applies to GI's conditionality, here applied to the audit's own structure.

The asymmetry runs through the fallbacks of §3 as well, in complementary directions. M1 carries the *discharge weight*: its failure suspends the discharge, relocating GI's conditionality to a single named step without striking anything. M2 carries the *refutation risk*: an invariant outside the holonomy algebra is the one fallback that could strike GI rather than suspend it. M3 carries the *consistency risk*: its failure alarms two arcs at once. A verifier's effort should be allocated accordingly — M1's construction step first, hardest, and with the as-transport / from-transport-constraints distinction in hand.

A useful way to hold the result is therefore this: the paper does not argue that two independently discovered circles happen to coincide. It argues that one of those circles may never have existed as an independent object. The conventional reading begins with an electromagnetic gauge phase and asks why it should be compact. The present reading begins with substrate transport, derives compactness, and audits the Maxwell construction; if that construction is genuinely transport all the way down, the electromagnetic phase is not a second circle waiting to be identified — it is the transport circle described in electromagnetic language, and the charge paper becomes a consequence of a single compact phase structure rather than a correspondence between two. [Calibration of the imports; no marker changes.]

4. The Exhibition — The Maxwell Connection as Substrate Transport

The exhibition establishes (GI-a), and under M1 it is short — deliberately so. The length of this section is itself a finding: if M1 holds as stated, there is almost nothing to exhibit, because the construction has already done the exhibiting. The work of this section is to say exactly what the identity consists in, at what grain it holds, and what it does and does not quantify over — because those distinctions carry §§5–7.

The two maps. Write $\mathcal{P}$ for the admissible paths of the substrate and $h_T : \mathcal{P} \to \mathbb{T}$ for the substrate transport holonomy assignment — the map whose value space the U(1) paper derives and whose path-by-path determination the Assignment Paper addresses. Write h_M for the holonomy assignment of the Maxwell connection — the map the electromagnetic sector uses, whose closed-loop values are the content of electromagnetic interference and whose infinitesimal form is the field strength. GI-a is the statement:

$h_M = h_T$ — one map, not two maps that agree.

What M1 delivers. Under M1, the Maxwell connection's parallel transport along $\gamma \in \mathcal{P}$ *is* the substrate transport along γ: the electromagnetic sector introduces no transport of its own, and h_M is h_T by construction — the same assignment presented in electromagnetic dress, with the gauge potential A as its local logarithmic description and the curvature F as its infinitesimal loop form. The identity is therefore not a theorem proved by comparing two independently defined maps on a common domain; it is the recognition that the construction defines one map twice. This is the strongest form an identification can take, and it is the form (GI-c) will need: identity-by-construction leaves no room for a hidden relabeling between two circles, because there is one circle (§6).

The grain of the identity. The identity $h_M = h_T$ is quantified over paths but fixes no values: it says that *whatever* holonomy a path carries, both descriptions carry that one. It is therefore a statement about the *form* of the assignment — about which map the electromagnetic sector uses — and not about the assignment's content. No step of this section requires knowing, for any particular path, which element of $\mathbb{T}$ it carries. This observation is banked now and spent in §7: the exhibition is assignment-neutral, and the independence the charge paper proves in its §12 survives the discharge.

What the exhibition leaves open. Identity of holonomy assignments fixes the values $h(\gamma)$ as elements of some group, but it does not by itself settle *which group* — the same assignment of unit-modulus factors could be read as valued in $\mathbb{T}$ or lifted to $\mathbb{R}$, depending on whether accumulations differing by 2π count as distinct. That is (GI-b), and it is not a formality: it is the odometer question, the exact point at which conventional electromagnetism goes undetermined and the substrate derivation must supply what assumption conventionally supplies. §5 settles it. [Exhibition: Proven, conditional on M1 [Inherited]; the identity is by construction, the section's content being its exact statement and grain.]

5. Reading Off the Structure Group — Why $\mathbb{T}$ and Not $\mathbb{R}$

The structure group of a connection is fixed by one question: *which holonomies count as the same holonomy?* If the accumulated phase θ is physical — if a transport that winds once and a transport that does not wind are distinct invariant configurations even when their unit-modulus factors agree — the group is the line $\mathbb{R}$, phase is an odometer, and nothing in §6 of the charge paper goes through: $\rho(\theta + 2\pi) = \rho(\theta)$ is not forced, and every real coupling is admissible. If only the unit-modulus factor $e^{i\theta}$ is physical — if θ and $\theta + 2\pi$ are one point — the group is the circle $\mathbb{T}$, and the lattice follows. Conventional electromagnetism cannot answer the question from inside: with charge quantization assumed, all observables are functions of $e^{iq\theta}$ with $q \in \mathbb{Z}$, and these are blind to the difference between $\mathbb{T}$ and $\mathbb{R}$ — which is why the conventional literature must reach for monopoles or unification to settle what the formalism leaves open. The substrate answers it, in three steps, each consuming a named input.

Step one — the admissible object is the exponential. The U(1) paper's compactness-with-provenance result [Inherited, as fixed in the charge paper §2]: reversible substrate transport acts by unit-modulus complex factors — the admissible object is $e^{i\theta}$, not θ. The real number θ is a coordinate on the circle, not a physical magnitude; selecting a branch of it — deciding that *this* preimage in $\mathbb{R}$ rather than *that* one is the real value — is a distinction no admissible invariant witnesses, and the exact-periodicity result forbids it as admissible structure. On the substrate side, the structure group of transport is $\mathbb{T}$, and this is inherited, not argued here.

Step two — the Maxwell sector adds no winding record at the fiber. The identity of §4 transfers the substrate's answer to the electromagnetic side *unless* the Maxwell sector independently promotes the lift — unless some electromagnetic construction treats total winding as invariant content of the holonomy itself that the substrate does not carry. M3 is exactly the statement that it does not: the real-valued potential and its bare line integrals enter the Maxwell sector as descriptive intermediaries, and every invariant is a function of holonomies — including the composed-holonomy data of realized curvature records, which is where the sector's real-valued invariants legitimately live (step three) — while none is a function of the lift. M2's positive load in the argument is exactly here, and is stated in its own words once: M2 guarantees the survey is complete — there is no invariant content outside the holonomy algebra, so read, in which a winding record could hide — so that "no electromagnetic odometer at the fiber" exhausts the sector rather than only its surveyed parts. This completeness role is also why M2's fallback is the sole GI-striking one (§3): an invariant outside the algebra would not merely weaken the survey but open territory the reading never covered. With M2 and M3 verified, there is no electromagnetic odometer at the fiber; the clock face of the substrate is the clock face of electromagnetism.

Step three — an in-principle winding record does not count. One residual worry has the same shape as the amplification objection the charge paper's §5 meets, and it is met the same way. A winding number *can* be witnessed in particular realized circumstances: a history that carries a record of its own circuit — a counter that ticks at each revolution — realizes an invariant distinguishing θ from $\theta + 2\pi$ *for that history*. Does the in-principle availability of such records re-open the lift? It does not, and for the inherited reason: the Conservation Lemma and the closure machinery run on *realized* witnesses, not on witnesses constructible in principle — the U(1) paper's architecture line "no amplification witness at unrealized limits" is this discipline, applied there to the catalogue boundary, in the charge paper to the response, and here to the lift. A realized circuit-counter is additional realized structure — a record, with its own invariant content — and distinguishing histories by their records is ordinary physics, not a property of the holonomy catalogue. The counter is not hypothetical; electromagnetism supplies its standard instance, and naming it converts the hardest apparent counterexample into a confirmation. Magnetic flux through a realized spanning surface, $\Phi(S) = \iint S$ *F, is real-valued, unbounded, and unimpeachably admissible — and it is exactly a realized record: integrated curvature, the composed holonomy data of the surface's small loops, carried by the realized configuration rather than by any fiber element. Where a potential exists, Stokes gives* $\Phi(S) = \oint \partial S$ A, which is precisely why two circulations differing by 2π *can* be physically distinct — as field configurations, witnessed by the curvature record over the surface, never as fiber points; two connections whose fluxes differ by 2π are different physical situations that assign the *same* holonomy to the boundary loop. Real flux is therefore not the lift; it is the realized record the lift would need and the bare circulation lacks. A three-object taxonomy fixes the section's result in one breath: the holonomy $h(\gamma) \in \mathbb{T}$ is fiber-level and circle-valued; the flux $\Phi(S) \in \mathbb{R}$ is configuration-level and record-borne, inside M2's algebra; the bare circulation $\oint\_\gamma$ A absent a realized record is a coordinate on the circle, not a point of it. The taxonomy introduces no new physical assumption; it distinguishes three objects that travel under the single label "phase accumulation," and the distinction is the entire content of the resolution. The lift M3 forbids is promoting the third to invariant standing, or reading the second as fiber data. The structure group is a catalogue-level object: it answers which holonomies are the same *absent* further realized structure, and absent further realized structure, no admissible invariant separates θ from $\theta + 2\pi$. One external anchor corroborates the discipline rather than extending it: in conventional gauge theory the gauge-invariant observables are exactly the Wilson loops — the exponentiated holonomies — and accumulated winding appears nowhere in the observable algebra, while flux enters through curvature over surfaces, exactly the record-borne route [Imported-External]. The realized-witness discipline is therefore not doing idiosyncratic work at this site; it recovers, from below, the observable algebra conventional gauge theory writes down by hand — the kind of convergence the corpus records where it occurs. The flux case is itself such a convergence: the discipline's category of "realized record" is exactly the category flux occupies, so the hardest apparent counterexample to the compact reading is not an exception the section must survive but an instance the section's own machinery classifies. The group is $\mathbb{T}$.

The reading. With the three steps in place the structure group is read off, and the reading is one line:

Structure group of the Maxwell connection = group of admissible transport holonomy values = $\mathbb{T}$.

(GI-b) is established. The odometer question — the exact question conventional electromagnetism cannot answer and the entire conventional apparatus of monopoles and unification exists to answer for it — is settled from below: phase is a clock face because a branch of the angle is an unwitnessed distinction, and a finite world sustains no structure nothing in it could witness. [Proven, conditional on M1–M3 [Inherited], the inherited exact-periodicity and compactness results, and the realized-witness discipline [Inherited].]

6. The Identification Is Canonical — Identity, Not Isomorphism

A check that two structures are "both $\mathbb{T}$" establishes nothing: every circle is isomorphic to every other as a topological group, and an audit that terminated at isomorphism would leave GI exactly where it started, with the identification an unexamined choice. The discharge requires canonicality: the identification between the electromagnetic circle and the transport circle must be fixed by the construction, point for point, with no residual freedom that physics could hide in. This section establishes that the residual freedom is exactly one bit — orientation — and that the bit is physically empty.

The classification. The continuous homomorphisms $\mathbb{T} \to \mathbb{T}$ are exactly the integer power maps $z \mapsto z^k$, $k \in \mathbb{Z}$ [Imported-External — classical; the continuous characters of the circle]. Among these, the automorphisms — the candidates for an identification — are $k = +1$ (the identity) and $k = -1$ (conjugation, $z \mapsto \bar{z}$). Every other surjective endomorphism, $|k| \geq 2$, is non-injective: it collapses k distinct points of the source circle to one point of the target.

Why $|k| \geq 2$ is excluded. Suppose the electromagnetic circle were related to the transport circle by $z \mapsto z^k$ with $|k| \geq 2$. Then distinct transport holonomies — k-th roots of a common value — would present identically on the electromagnetic side: the electromagnetic phase would be strictly *coarser* than transport phase, losing admissible distinctions the substrate carries. That is not a relabeling; it is a genuine non-identity, a second circle after all, related to the first but not the first. M1 excludes it at the root: identity-by-construction means there are not two circles between which a power map could intervene — h_M *is* h_T, and the identification is the identity map because it is not a map between different things at all. The k-fold scenario could only arise in the construction-from-constraints fallback of §3, where an independently introduced electromagnetic phase variable might wind around transport phase; under M1 as stated, it cannot arise. This is the precise dividend of M1's strength, and the reason its verification obligation insists on *as transport* rather than *from transport constraints*.

The surviving bit. What identity-by-construction does not fix is orientation: the realization of the transport circle in $\mathbb{C}$ involves the choice of i over −i, and conjugation $z \mapsto \bar{z}$ implements the exchange. Under conjugation every representation label flips sign, $q \mapsto -q$: the carrier that wound q times now winds −q times. This is the charge-sign convention — which species is called positive — and it is physically empty in exactly the sense the corpus requires: no admissible invariant distinguishes the two orientations, since every invariant is a function of holonomies and their conjugates symmetrically. The freedom is real, conventional, and harmless; it is the same freedom conventional electromagnetism carries in the sign of e, and it is the *only* freedom the identification retains.

Periodicity, transferred verbatim. With the identification canonical up to orientation, the 2π-periodicity of the electromagnetic phase is the *same* identification as the 2π-periodicity of the transport phase — one point of one circle, not two matched points of two circles. The closure condition the charge paper's §6 consumes — $\rho(\theta + 2\pi) = \rho(\theta)$, whence $e^{2\pi iq} = 1$, whence $q \in \mathbb{Z}$ — therefore applies to the electromagnetic phase verbatim, with no rescaling, no hidden winding, and no relabeled generator. The lattice the charge paper derives on the transport circle is the lattice on the electromagnetic circle, because they are one circle. (GI-c) is established. [Proven, conditional on M1 [Inherited]; the classification of endomorphisms is [Imported-External]; the orientation freedom is recorded as convention, no marker.]

7. Independence from the Assignment Node, Preserved

The charge paper's quiet structural claim — that the charge chain does not wait on the assignment node — was proved there by an input-by-input audit and re-verified after each repair. The discharge must not silently spend that independence, and this section runs the audit once more over the present paper's inputs.

The inputs consumed above are: M1, the identity of the two holonomy assignments (§4); M2, holonomy exhaustion (§3, consumed in §5 step two as the completeness guarantee for the invariant survey); M3, no independent lift (§5); the inherited catalogue results — compactness, exact periodicity, the Conservation Lemma, the realized-witness discipline (§5); and the classification of continuous endomorphisms of $\mathbb{T}$ (§6).

None consumes the assignment. M1 is an identity of maps quantified over paths: it says h_M and h_T are one map, while fixing the value of that map at no path — two rival assignments sharing the catalogue would each satisfy M1 by carrying their own values identically in both descriptions. M2 and M3 constrain which *kinds* of invariants exist, not which paths carry which values. The catalogue results are catalogue-level at source. The endomorphism classification is mathematics, consuming nothing physical. The discharge is therefore assignment-neutral in the same exact sense as the charge paper's chain: a failure of Transport-Completeness in the Assignment Paper would leave the present reading, and the discharged GI, standing untouched. The complementarity recorded in the charge paper's §12.1 — the two arcs bearing on GI from opposite directions — is superseded in one respect only: the strengthened, Assignment-conditional form of the GI refuter is absorbed by the discharge (§8), since a premise that has been verified no longer needs its burden-shift reinforced. [Proven — the preserved independence — given the stated grain of every consumed input.]

8. Consequences — The Cascade Through the Charge Paper

The discharge propagates through the corpus by marker arithmetic, and the arithmetic is recorded exactly.

GI. From [Conditional — named premise of the charge paper] to [Proven, conditional on M1–M3 [Inherited] and the inherited compact-phase result]. The premise is not eliminated; it is *relocated* — converted from a free assumption of the charge paper into a dependence on the Maxwell arc at its source statuses. Honesty requires the relocation to be stated as one: the corpus's exposure is unchanged in extent and changed in kind, from an unverified bridge between two arcs to the ordinary inherited exposure of the arcs themselves.

The charge lattice. From [Proven, conditional on GI] to [Proven, conditional on the corpus's inherited inputs — the U(1) arc and the Maxwell arc at their source statuses]. The charge paper's chain now runs from the programme's foundations to the integer lattice with no free premise in it. Its text requires no revision: the conditioning was carried in every marker precisely so that this discharge could land without touching the paper.

The refuters. The charge paper's §13 load-bearing refuter — exhibit an admissible electromagnetic phase invariant not representable as transport holonomy — remains the operative falsification form, but its target shifts upstream: such an invariant would now refute M2 at source rather than an open premise. The strengthened Assignment-conditional refuter is absorbed: the inversion it existed to close (the challenger's option of declaring the transport phase the idle auxiliary) is closed unconditionally by the discharge, since a phase that the electromagnetic sector is *constructed as* cannot be idle in it. The empirical channel consolidates: an observed off-lattice charge now strikes, jointly and at once, the U(1) arc's compactness, the Maxwell arc's M1–M3, and the charge paper's representation chain — one lever, three arcs, which is the correct shape for a programme whose parts are meant to be load-bearing rather than modular decorations.

The programme diagram. The branch this paper completes:

```
Admissibility of the measure
        ⟹ (companion) finite holonomy impossible ⟹ H infinite

H infinite + subgroup ⟹ H dense in 𝕋

Finite distinguishability + no amplification witness at unrealized limits
        ⟹ Conservation Lemma ⟹ no boundary stands ⟹ H closed

Dense + closed ⟹ admissible phase structure = 𝕋
Reversible (ℂ-realized), single-Fold ⟹ exactly U(1)
         │
         ├── U(1) ⟹ Born obligation discharged
         │
         ├── U(1) + GI ⟹ charge lattice
         │        └── GI discharged: Maxwell connection exhibited as
         │            substrate transport; structure group read as 𝕋;
         │            identification canonical up to orientation
         │            [Proven, conditional on M1–M3 — this paper]
         │
```

```
    └── U(1) + assignment L ↦ h(L) D-determined up to gauge
              [Conditional — Assignment Paper, on Transport-Completeness of
D]
              ⟹ support supervenience ⟹ dyadic loading
```

What does not cascade. The unit of charge remains unfixed (the fine-structure arc's business, exactly as the charge paper's §10 divides it); the occupancy pattern and the e/3 generator remain the multi-Fold arc's burden (charge paper §9); and the assignment node remains at the Assignment Paper's [Conditional], untouched (§7). The discharge closes one obligation and opens none.

8.1 Why the Discharge Matters Beyond Charge

The discharge does more than strengthen one result, and the larger significance is recorded once, as calibration. The charge lattice is the first Standard-Model-facing result in the corpus whose final free premise is removed by an audit rather than by a new postulate — and that distinction changes the programme's explanatory direction. Conventionally, charge quantization is explained by *adding* structure: monopoles, grand-unified embeddings, enlarged symmetry, each purchasing the lattice with physics never observed. Here the direction reverses: a premise is removed, no structure is added, and the lattice survives the removal. The honest qualification of §8 stands alongside this and is not in tension with it — the corpus's exposure relocates into the Maxwell arc rather than vanishing — because relocation into already-load-bearing structure is exactly what distinguishes an audit from a postulate: nothing is believed at the end of the audit that was not already believed for other reasons.

Why the monopole question changes status. One corollary should be banked rather than left for readers to notice, and its framing matters as much as its content. In conventional physics magnetic monopoles are introduced *to explain* charge quantization — the logic runs *monopole existence ⟹ charge quantization*, and the two questions travel together as if inseparable. The present route reverses the dependence: quantization is derived with no monopole premise, compactness inherited from the transport phase itself, and the monopole question thereby changes status. It is no longer part of the explanation of charge quantization; it becomes a downstream question about what the construction permits. The discharge therefore separates two questions the literature usually treats as one — *why is charge quantized?* and *do magnetic monopoles exist?* — and within this framework the first is addressed before, and independently of, the second.

The downstream question has exact mathematical form. A compact structure group makes available what a non-compact one cannot: nontrivial bundle topology — sectors of nonzero first Chern class, the mathematical home of monopoles. Does the substrate construction of the Maxwell sector permit or forbid such sectors? Nothing here answers it. The question is recorded as a named successor obligation at [Open], and it is a question only available to ask in substrate terms because compactness is derived from a construction fine-grained enough to interrogate; either answer would be a finding, since a forbidden monopole sector would be a prediction — no magnetic monopoles of the bundle (Dirac) type — and a permitted one would re-derive Dirac's condition from below.

Whether the route succeeds turns on the verification of M1–M3, and nothing here pre-empts that. But conditional on the slots being fixed as stated, one of the Standard Model's oldest unexplained regularities traces to a phase structure already required for quantum interference — to the circle the Born rule runs on — rather than to new particles or new symmetries. The explanation of charge discreteness costs the corpus nothing it has not already spent. [Calibration; no marker.]

9. What Would Refute This

The audit's refuters divide into refuters of the reading and failures of the imports; both are named, with the fallbacks of §3 cross-referenced rather than restated.

Against M1 — a non-transport component. Exhibit a component of the Maxwell connection not constructed as substrate transport: a step in the source construction at which the electromagnetic phase enters as an independent degree of freedom rather than as transport phase presented electromagnetically. Effect: the identity of §4 fails at that step; GI reverts to [Conditional] with the gap exactly located; §§5–6 survive as conditional machinery (fallback, §3). Standing: under M1 as stated no such component exists; the verifier's obligation is to confirm the statement against the source.

Against M3 — a realized winding witness at the fiber. Exhibit a realized admissible invariant that distinguishes holonomies differing by 2π absent further realized record structure — an electromagnetic odometer at the fiber. Effect: the structure group reverts to $\mathbb{R}$, refuting not only the present reading but the U(1) paper's exact-periodicity result upstream; this is the inconsistency alarm of §3, striking two arcs at once, and it is the strongest single refuter available against the corpus's phase programme. Standing: the in-principle availability of circuit-counters does not qualify, by the realized-witness discipline (§5, step three); magnetic flux through a realized spanning surface does not qualify either — it *is* the realized record, configuration-level curvature data distinguishing field configurations, not fiber points (§5, step three). The witness demanded must separate fiber points themselves: two transports, identical in holonomy and identical in every realized record, physically distinguished nonetheless. The burden is that witness, at catalogue level.

Against the canonicality (§6). Exhibit a step of the source construction at which the electromagnetic phase variable winds k times ($|k| \geq 2$) around transport phase — a power-map relation surviving in the construction. Effect: the electromagnetic phase is coarser than transport phase; (GI-c) fails; the lattice survives as a lattice but GI as identity fails. Standing: excluded by M1's identity-by-construction; this refuter is live only in the construction-from-constraints fallback, where it becomes the first thing to check.

Against the verification itself. The audit is complete only when the three slots of §3 are fixed against the source papers. An unfixed slot is not a discharged premise: until citation, the present paper's central results stand at [Proven, conditional on M1–M3 *as stated*], and GI's effective corpus status is the weaker of that and its prior [Conditional]. One asymmetry within this standing: a construction-level M3 violation discovered during slot verification does not leave the paper at ordinary unverified standing — it triggers the inconsistency alarm of §3, striking the U(1) arc as well; the M3 slot can fail upward into a corpus problem in a way the other two slots cannot. The completion condition is administrative but absolute, and it is recorded here so that no reader mistakes a stated import for a verified one.

Empirically. The consolidated channel of §8: an off-lattice charge — any irrational charge ratio, any continuously tunable charge — strikes the conjunction of the three arcs. Standing: Millikan-type searches and the part-in-10^{21} electron–proton balance sit entirely on the lattice, exactly as the charge paper records.

Inherited exposure. Everything above is additionally conditional on the U(1) arc's inputs at source — no-pre-individuation, reversibility, finite-holonomy-impossibility, single-Fold binarity — and on the Maxwell arc's own foundations at theirs. The exposure is the corpus's, shared, and not re-audited here.

10. What the Paper Establishes

Established (conditionally, with the conditions named):

- The discharge theorem: M1 ∧ M2 ∧ M3, with the inherited compact-phase result, entail GI in full strength — identity of holonomy (§4), compact structure group (§5), canonical identification up to orientation (§6). [Proven, given the imports as stated.]
- The structure-group reading: the group of the Maxwell connection is $\mathbb{T}$, not $\mathbb{R}$ — the odometer question settled from the substrate side, with the in-principle-winding loophole closed by the realized-witness discipline (§5). [Proven, conditional on M1–M3 and inherited inputs.]
- The flux resolution and the three-object taxonomy: the holonomy $h(\gamma) \in \mathbb{T}$, fiber-level and circle-valued; the flux $\Phi(S) \in \mathbb{R}$ through a realized spanning surface, configuration-level and record-borne — integrated curvature, composed holonomy data, inside M2's algebra; the bare circulation $\oint_\gamma A$ absent a record, a coordinate on the circle and not a point of it. Real-valued electromagnetic flux is thereby reconciled with the compact structure group — it distinguishes field configurations, never fiber points — and the hardest apparent counterexample to the compact reading is classified by the section's own machinery rather than survived by it (§5, step three). [Proven, conditional on M2, M3, and inherited inputs.]
- The canonicality result: the identification's residual freedom is exactly orientation, $z \mapsto \bar{z}$, implementing $q \mapsto -q$ — the charge-sign convention, physically empty; power-map relabelings $|k| \geq 2$ excluded by identity-by-construction (§6). [Proven, conditional on M1; endomorphism classification [Imported-External].]
- The exact import inventory: M1–M3 stated in the form the argument requires, each with citation slot, verification obligation, and named fallback — including the M2 fallback's conversion of the audit into a refutation procedure and the M3 fallback's inconsistency alarm (§3).
- The import asymmetry: M1 determines whether there are one or two phase structures and so carries the discharge weight; M2 carries the sole GI-striking fallback; M3 the cross-arc consistency alarm — an allocation guide for the verifier, with the as-transport / from-transport-constraints distinction as the first and hardest check, and identity-by-

construction recognized as categorically stronger than correspondence: the identification questions dissolve rather than require solving (§3.1). [Calibration.]

- Assignment-neutrality of the discharge: every consumed input is catalogue-level or form-level; the charge chain's independence from the assignment node survives, and the Assignment-conditional strengthened refuter is absorbed rather than violated (§7, §8). [Proven, given the stated grain of the inputs.]
- The cascade: GI to [Proven, conditional on M1–M3]; the charge lattice's conditionality relocated from a free premise into the Maxwell arc's source statuses; the empirical channel consolidated into one lever bearing on three arcs (§8).
- The programme-level calibration: the first Standard-Model-facing result whose final free premise is removed by audit rather than by postulate — structure removed, not added, with exposure relocated into already-load-bearing arcs (§8.1). [Calibration; no marker.]

Not established (open, out of scope, or pending):

- the verification of M1–M3 against the Maxwell-sector papers — the audit's completion condition; the slots of §3 are to be fixed at source, and until they are, the discharge stands at the conditional form recorded in §9. The M3 slot routes one-hop through the companion paper *From Real Circulation to Compact Holonomy*, which performs the construction-level check and itself consumes M2 — fixing an ordering constraint: the M2 slot first or jointly;
- monopole sectors — whether the substrate construction of the Maxwell sector permits or forbids nontrivial bundle topology (nonzero first Chern class), made mathematically available by the compact structure group; the discharge separates *why is charge quantized?* from *do monopoles exist?*, answering the first independently of the second and converting the second from a premise of the conventional route into a downstream question with either answer a finding (§8.1) — [Open], named successor obligation;
- the unit of charge (fine-structure arc), the occupancy pattern and the e/3 generator (multi-Fold arc), and the assignment question (Assignment Paper, at its status) — all exactly as divided in the charge paper, none advanced here;
- all inherited inputs of both arcs, at their source statuses.

The honest summary: this paper proves that three exactly stated properties of the Maxwell-sector construction suffice to discharge the charge paper's one free premise, settles the two subtleties — the odometer question and the canonicality question — on which a careless discharge would silently founder, and reduces the remaining work to fixing three citations against the source papers, with the failure mode of each fixed in advance. The corpus's conditional count goes down by one on the day the slots are filled, and the paper says precisely what happens on any day they cannot be.

11. Conclusion

The charge paper ends with one dial unverified: that the circle electromagnetism turns on is the circle the substrate derives. The verification, it turns out, is not a matching of two dials but the opening of a casing: the Maxwell sector, built as transport, contains no second dial to match — one spindle, presented twice, once as the phase of interference and once as the phase of charge.

What remains after the exhibition is the two questions a hasty audit would skip, and they are the questions this paper exists to settle. The dial is a clock face and not an odometer, because a branch of the angle is a distinction nothing admissible witnesses, and an unrealized circuit-counter does not count. And the two presentations agree point for point and not merely in shape, because identity-by-construction leaves no room for a hidden winding — the only surviving freedom is which direction to call positive, and nature does not care.

The significance of the paper is therefore calibrated exactly. It is not that a new physical consequence is derived — the charge paper already supplied the consequence. It is that the final free bridge in that derivation is removed, up to three named citations, by showing that the electromagnetic phase and the transport phase are not merely similar structures audited into correspondence but one structure viewed through two descriptions. A bridge between identical banks is not a bridge, and the paper's work is to establish the identity rather than to engineer the crossing.

Granting the three cited results of the Maxwell arc, the chain closes. A finite world, composing without bound, sustains no structure nothing in it could ever witness; its phase is therefore a circle; electromagnetism, made of its transport, turns on that circle because it is made of it; and a gear on a circle meshes only at whole numbers. The lattice of charge was never a brute fact and never an extra law. It is the sound of one gear, on one circle, closing its loop.