

Charge Quantization from Compact Phase

The Lattice from the Circle — Integer Charge as a Conditional Consequence of the Derived U(1)

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The question

One of the strangest plain facts about the world is that electric charge comes in fixed amounts. Every electron carries exactly the same charge as every other electron, to the limits of the most precise measurements ever made. Protons carry exactly the opposite. Quarks carry exactly one-third and two-thirds of that unit — never 0.31, never 0.6667001. Nature appears to use a strict ladder of charge values and nothing in between, and the Standard Model, for all its success, does not explain why. It writes the ladder down; it does not derive it. The conventional explanations on offer — grand unified theories, magnetic monopoles — each work by adding new physics that has never been observed.

What this paper argues

One paper in this programme — *Why Finite Distinguishability Forces Continuous U(1) Phase*, referred to throughout as the U(1) paper — derived something unusual: that the phase of quantum theory — picture it as a dial whose hand sweeps as physical histories unfold — is forced to be a complete, unbroken circle. Not a line that runs on forever, and not a set of discrete clicks, but a circle, closed and continuous, with positions that repeat exactly after one full turn. That result was forced from the substrate's own structure, not assumed.

This paper draws out a consequence. If phase is a circle, then anything that responds to phase must respond in a way that comes back to where it started after one full turn of the dial. Think of a gear meshed to the dial: when the dial completes one revolution, the gear must land exactly on a position it has already visited, or the meshing is inconsistent. The only consistent gear ratios are whole numbers — the gear may turn once, twice, three times per revolution of the dial, but never 1.37 times, because 1.37 turns does not return the gear to a visited position when the dial closes its loop.

Electric charge, in modern physics, *is* that gear ratio: it measures how strongly a particle's own internal dial turns in response to the phase. So if phase is a circle, charge is forced onto a ladder of whole-number multiples of a basic unit. The discreteness of charge stops being a brute fact and becomes a consequence of the shape of phase.

The honest condition

One assumption stands between the previous paper's result and this conclusion, and this paper names it rather than hiding it. The circle derived previously is the phase of quantum interference — the phase behind the Born rule. The phase that charge responds to is the phase of electromagnetism. Standard physics treats these as closely related but distinct structures. For the argument here to go through, the electromagnetic phase must be one and the same circle as the substrate's transport phase — one dial, not two. There are good structural reasons inside the programme to expect this — the substrate's phase is already path-dependent in exactly the way electromagnetic phase is — but the identification is carried here as a named premise with a named way to refute it, not as a proven step. If it holds, charge quantization follows from the programme's foundations with nothing added. If it fails, this paper says precisely where and how.

What the paper does not claim

The argument fixes the *ladder* — that charges are whole-number multiples of a unit. It does not fix the *size* of the unit (that is the business of the programme's fine-structure work), and it does not by itself produce the specific fractions $\pm 1/3$ and $\pm 2/3$ carried by quarks, which arise from how electromagnetism sits inside the larger structure of the Standard Model. Discreteness is derived here; the particular rungs occupied are downstream work.

Abstract

Electric charge quantization — the empirical fact that observed charges occupy a discrete lattice — is unexplained within the Standard Model, where charge assignments are inputs. The conventional derivations (grand unification, Dirac monopoles, anomaly cancellation) each purchase quantization with additional physics or additional consistency conditions. A more economical observation is standard but rarely usable: charge quantization follows already if the gauge group of electromagnetism is the compact $U(1)$ rather than the non-compact \mathbb{R} , since the continuous one-dimensional unitary representations of $U(1)$ are exactly $\theta \mapsto e^{iq\theta}$ with $q \in \mathbb{Z}$. The observation is rarely usable because conventionally the compactness is itself an assumption — precisely the assumption that grand unification and monopoles exist to justify.

The $U(1)$ paper of this programme (*Why Finite Distinguishability Forces Continuous $U(1)$ Phase*) derives the compactness: the substrate's transport phase is forced to be compact, connected, one-dimensional, and abelian — exactly $U(1)$ — by reversibility, unbounded composition, and no-pre-individuation, at the level of the holonomy catalogue [Inherited]. The present paper converts that result into charge quantization in four steps, with statuses tracked. (i) **The gauge identification** — that the electromagnetic gauge phase is the substrate transport phase, equivalently that the structure group of the Maxwell connection is the derived compact \mathbb{T} rather than \mathbb{R} — is isolated as the paper's named premise GI, motivated and audited (§3, §3.2): four structural coincidences support it (shared holonomy form, shared carrier, the Maxwell sector's transport character, and the no-pre-individuation burden against an idle second circle), the last shifting the burden onto any second-circle alternative to exhibit independent admissible

invariant content; GI is nonetheless carried as [Conditional] with its sharpened refuter (§13). (ii) The carrier is the candidate-layer structure transported along admissible paths; that its response factors through accumulated holonomy at all — **response supervenience on holonomy**, the well-definedness premise conventional treatments cross silently — is forced by the Admissibility Conservation Lemma, and the homomorphism and unitarity of the response then follow from composition and unit-modulus reversibility [Proven, given inherited inputs] (§4). (iii) The **continuity of the response** — normally a suppressed textbook assumption, and a genuine loophole, since discontinuous characters of $U(1)$ exist — is forced in two layers: admissible continuity (no response gap without an invariant change) directly from the Conservation Lemma, and topological continuity derived from it via the homomorphism property, with the classical fact that every measurable character of $U(1)$ is continuous as backstop — the discontinuous characters being exactly the non-measurable, choice-constructed ones, specifiable by no admissible construction [Proven, conditional on inherited inputs] (§5). (iv) Compactness then makes the classification unconditional: $\rho(\theta + 2\pi) = \rho(\theta)$ forces $e^{i2\pi q} = 1$, hence $q \in \mathbb{Z}$ — the charge lattice [Proven, conditional on GI] (§6–7).

The contribution claimed is calibrated explicitly (§11): the representation-theoretic step is textbook; what the programme supplies is the *premise* — derived compactness — that the textbook step has always lacked, replacing the unobserved structures (GUT embeddings, monopoles) that conventionally supply it. The paper additionally establishes: that the empirical discreteness of charge runs the inference backwards as independent observational support for the compact-phase result, with the caution that this support cannot discriminate compact $U(1)$ from its dense proper subgroups (§8); that fractional quark charges are an embedding question downstream of the derived lattice, not a counterexample to it — with the lattice's one crisp standalone prediction stated as such: no charge will be observed off integer multiples of $e/3$ (§9); that the lattice spacing — the value of the elementary charge — is not fixed here and is the business of the programme's coupling-constant work (§10); and that the entire result hangs off the catalogue half of the $U(1)$ paper's theorem and is therefore **independent of the assignment question** — a node standing at [Conditional] at the Assignment Paper's status — the first piece of Standard-Model-facing physics in the programme that does not wait on it (§12, §12.1).

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 Question and the Route
 2. The Inherited Phase Result
 3. The Gauge Identification — The Paper's Named Premise
 - 3.1 Why the Derived Phase Is Not a Free Circle
 - 3.2 The Gauge Identification Audit — Why GI Is Not an Arbitrary Bridge
 4. What Carries the Representation, and Why It Is a Representation
 5. Continuity of Response, Forced — Closing the Discontinuous-Character Loophole
 6. The Compactness Constraint — Integer Labels
 7. Charge as the Representation Label
 8. Why Non-Compact Phase Fails — and the Inference Run Backwards
 9. Fractional Charges — Embedding, Not Exception
 10. The Unit of Charge — What This Paper Does Not Fix
 11. Comparison with the Conventional Routes
 12. Position in the Programme — Independence from the Assignment Node
 - 12.1 Why the Charge Result and the Assignment Result Are Complementary
 13. What Would Refute This
 14. What the Paper Establishes
 15. Conclusion
-

1. Introduction — The Question and the Route

Charge quantization is among the oldest unresolved structural questions in physics. Observed electric charges occupy a discrete lattice: integer multiples of the proton charge for leptons and hadrons, with quark constituents at thirds. No continuously varying charge has ever been observed, and the equality of charge magnitudes across particle species — electron against proton — is confirmed to better than one part in 10^{21} . Quantum electrodynamics takes the charge assignments as inputs. The question *why the lattice* is not answered inside the Standard Model.

Three conventional routes exist, and each purchases the answer with something extra. Grand Unified Theories obtain quantization by embedding electromagnetic $U(1)$ inside a larger compact simple group, whose representation theory discretizes the charges — at the price of new gauge bosons, new symmetry-breaking structure, and proton decay, none observed. Dirac's monopole argument obtains the quantization condition $qg = 2\pi n$ from the consistency of quantum mechanics in the presence of a single magnetic monopole — at the price of the monopole, not observed. Anomaly cancellation constrains the hypercharge assignments of the Standard Model fermions tightly enough to force charge ratios among them — the strongest conventional argument, but one that fixes ratios within the observed particle content rather than explaining why charge is lattice-valued at all (§11).

There is a fourth observation, standard in the textbooks and almost never usable: quantization follows already if the gauge group of electromagnetism is the compact circle $U(1)$ rather than the

non-compact line \mathbb{R} . Compactness alone discretizes the one-dimensional representations. The observation is rarely usable because the compactness is, conventionally, an assumption — indeed it is exactly the assumption that grand unification and monopoles exist to justify. One assumes compact $U(1)$ and gets quantization, or assumes quantization and infers compactness; the circle of explanation does not close.

The $U(1)$ paper changes the standing of that fourth route. *Why Finite Distinguishability Forces Continuous $U(1)$ Phase* derives the compactness: the substrate's transport phase is forced — by reversibility, unbounded composition, and the no-pre-individuation discipline — to be exactly $U(1)$, compact, connected, one-dimensional, and abelian, at the level of the holonomy catalogue [Inherited]. If that derivation stands, the premise the textbook route has always lacked is supplied from below, and charge quantization follows with no unobserved structures added.

One conditional stands between the inherited result and the conclusion, and the route through this paper is organized around naming it rather than crossing it silently. The $U(1)$ derived in that paper is the candidate-layer transport phase — the phase of the Double Square kernel, the phase whose continuity makes the Born measure quadratic. The $U(1)$ of electromagnetism is a gauge phase. These are distinct structures in conventional physics. The identification of the two is the paper's named premise, stated, motivated, and given its refuter in §3, and carried conditionally through everything downstream. The remaining steps are then made at the programme's standard: the representation premise is grounded rather than supposed (§4); the continuity of the representation — a genuine mathematical loophole that the textbooks suppress by appeal to physical reasonableness — is forced by the programme's own Conservation Lemma rather than assumed (§5); and the final classification is then unconditional (§6).

2. The Inherited Phase Result

The $U(1)$ paper establishes, in the structural form of its §7.2, the following [Inherited, with the conditional status carried at source]:

The admissible phase structure of substrate transport is the full circle of unit-modulus phase factors,

$$\mathbb{T} = \{ z \in \mathbb{C} : |z| = 1 \} = U(1),$$

compact, connected, one-dimensional, and abelian. The chain that forces it: admissible holonomies form a subgroup of \mathbb{T} (reversibility, composition); the subgroup is infinite (the companion's finite-holonomy-impossibility) and hence dense; no admissible mechanism can maintain a catalogue boundary (the Boundary Maintenance trichotomy; the Admissibility Conservation Lemma), so the structure is closed; dense and closed in \mathbb{T} is \mathbb{T} ; and the ambient group — realized in \mathbb{C} from the outset, which excludes the solenoidal alternatives — is exactly $U(1)$.

Three features of the inherited result do work below and are flagged here.

Compactness with provenance. The phase group is the circle, not the line, *because* reversible transport acts by unit-modulus complex factors: the admissible object is $e^{i\theta}$, not θ . Compactness is not a regularity assumption but a consequence of reversibility. This is the feature that the entire present paper converts into the charge lattice.

Exact periodicity. On $U(1)$, phases differing by 2π are not approximately but exactly identical: there is one point of the circle, not two nearby points. No admissible structure selects a branch of the angle — branch selection is an unwitnessed distinction, forbidden. The identification $\theta \equiv \theta + 2\pi$ is therefore not a convention but admissible structure, and it is the identification that §6 turns into integrality.

The Admissibility Conservation Lemma. [Inherited] Physical admissibility may change only where some admissible invariant changes. Derived at source from no-pre-individuation: an admissibility difference unaccompanied by any invariant difference would itself be an unwitnessed distinction. The Lemma closed the catalogue there; here it closes a different gap — the continuity of the charge representation (§5).

The inherited result is a statement about the catalogue — the value space of the phase — and not about the assignment of holonomy to individual paths, which the $U(1)$ paper carries at [Open] and the Assignment Paper addresses, concluding D-determination up to gauge conditionally on its named Transport-Completeness premise (§12, §12.1). §12 establishes that nothing in the present paper consumes the assignment, under either status; the charge result hangs entirely off the catalogue.

3. The Gauge Identification — The Paper's Named Premise

Here is the step that conventional presentations of "compactness implies quantization" do not need and this programme does. The compact $U(1)$ of §2 is the substrate transport phase: the phase of admissible candidate-layer transport, the phase of interference, the phase of the Born measure. Electric charge, in gauge theory, couples to the gauge phase of electromagnetism: a local phase with its own connection, whose curvature is the electromagnetic field. In conventional physics these are intimately related — the gauge phase acts on the same wavefunction whose global phase is the interference phase — but they are not the same object, and a derivation of the compactness of one does not automatically constrain the other.

The premise is therefore stated as what it is.

Gauge Identification (GI) — *named premise of this paper* [Conditional]

The gauge phase of electromagnetism is the substrate transport phase: the phase to which charge couples ranges over the same compact circle that §2 derives, with the same exact 2π -periodicity. Stated in gauge-theoretic terms, GI is precisely the claim that the structure group of the Maxwell connection is the derived compact \mathbb{T} rather than \mathbb{R} — the gauge phase *is* transport holonomy, not a second circle and not a line.

A weaker-sounding formulation suggests itself and is rejected deliberately: that the gauge phase might be a distinct structure that merely *inherits* compactness from the transport phase. Inheritance-without-identity would require a mechanism — a construction assembling the gauge phase from transport holonomy, say as the structure group of a connection built from substrate transport — and nothing in this paper supplies one. A weaker-sounding premise that is actually unsupported is worse, by the programme's standards, than a stronger premise honestly carried. GI is the identification claim, cleanly, and everything downstream is conditioned on it as such.

Everything downstream of this section is conditional on GI, and the conditioning is carried in every marker. The structure of the paper's claim is therefore exact: *derived compactness (inherited) + GI (named, conditional) \Rightarrow charge lattice (proven given both)*. A reader who rejects GI keeps the inherited result and loses only this paper; a reader who grants it gets charge quantization from the programme's foundations with nothing else added.

Three considerations motivate GI from inside the programme — the structural form of the substrate phase (path-dependent holonomy, the signature of a connection rather than a global symmetry), the single-Fold economy (the substrate supplies one independent compact phase, not two), and the transport character of the Maxwell sector — and §3.2 develops them, expanded to four structural coincidences, as the single authoritative treatment, extracting the refuter's exact form. The Maxwell-sector consideration carries the concrete successor obligation, stated once here in structure-group form: exhibit the Maxwell-sector phase as transport holonomy and read off its structure group — nothing more, nothing vaguer. The present paper carries GI at [Conditional] pending that audit; §13 names its refuter.

3.1 Why the Derived Phase Is Not a Free Circle

The Assignment Paper sharpens the standing of the transport phase in a way that bears directly on GI, and the bearing should be recorded — with its conditionality. That paper addresses the node the U(1) paper carries at [Open]: whether the holonomy assignment $L \mapsto h(L)$, and not only its value space, is D-determined. Its conclusion is that the physically meaningful assignment is determined by D up to gauge — a difference in closed-loop holonomy is either witnessed by some admissible comparison, hence recorded in D, or unwitnessable by every admissible comparison, hence not physical structure at all — carried at its source status: [Conditional] on the Transport-Completeness of D, that paper's named and undischarged premise, with the holonomy principle (gauge-invariant content is closed-loop content) [Imported-External].

If that conclusion stands, the transport phase is not merely a compact catalogue of available values. Its gauge-invariant content is fixed by admissible distinguishability itself: the phase is a physically operative component of admissible comparison — carrying content through interference, support, and transport — not a freely assignable auxiliary structure. The question GI poses then changes character. It is no longer whether electromagnetism should use one among many available phase circles; on the programme's account there is already a distinguished, D-determined compact phase structure doing physical work. A second independent electromagnetic phase would therefore require not merely an additional circle but an additional *physically effective* phase structure beyond the one admissibility already forces — and the burden on that

alternative rises accordingly: it must exhibit admissible invariant content not already represented by transport holonomy, or fall as surplus structure rather than stand as an explanatory competitor.

Two honesty notes bound the strengthening. First, it is conditional twice over: on the Assignment Paper's conclusion, which is itself conditional on Transport-Completeness. Second — and structurally important — GI's standing in *this* paper does not depend on it. The charge result consumes only the catalogue (§12), GI's marker remains [Conditional] with or without the Assignment Paper, and this subsection strengthens the motivation, not the logic. If Transport-Completeness fell, this subsection would be deleted and nothing downstream of §3 would change. [Strengthened motivation; Conditional on the Assignment Paper at source status.]

3.2 The Gauge Identification Audit — Why GI Is Not an Arbitrary Bridge

GI is the paper's one new conditional, and a conditional that everything downstream consumes deserves more than motivation: it deserves an audit. The question the audit puts is not whether both structures are called $U(1)$ — that would be too weak, since names are free — but whether the electromagnetic gauge phase and the substrate transport phase perform the same physical role. Four structural coincidences say they do, and the fourth changes the premise's logical posture.

First — both are holonomy phases. The substrate phase is accumulated by transport along admissible paths, with closed-loop phase as the invariant residue. Electromagnetic phase has exactly the same form: what is physically meaningful is never an absolute local phase label but the phase accumulated around paths — electromagnetic interference, and Aharonov–Bohm-type effects, in which a charged amplitude acquires a measurable phase determined by the loop it encloses, are loop-holonomy phenomena outright. Nor is the force sector an exception: the electromagnetic field strength is the curvature of the connection — the infinitesimal form of loop holonomy — so even the force phenomenology of electromagnetism is holonomy phenomenology. This is not a superficial analogy between two structures that happen to share a name. It is the same mathematical role: path-dependent phase with closed-loop holonomy as the observable content. A framework in which phase is born path-dependent does not face the usual gap between global and gauged phase: the identification GI must make is of one connection with another, not the promotion of a global symmetry to a local one — a narrower gap than the conventional route must cross.

Second — both act on the same carrier. The phase of the Born-rule chain is the phase governing interference of candidate-layer amplitudes. The electromagnetic gauge phase is also the phase through which charged quantum amplitudes transform. If these were two independent phase structures, the framework would owe an account of why matter carries two independent $U(1)$ responses occupying the same amplitude layer — and the single-Fold sector supplies one independent compact phase, not two [Inherited]. The carrier of §4 would have to respond to both circles at once, with no admissible structure to keep the responses apart.

Third — the Maxwell sector is already a transport sector. The programme's Maxwell admissibility results construct electromagnetism from admissible transport constraints rather than as an added force. If the electromagnetic connection arises from admissible transport, its

structure group should be *read* from the admissible transport phase, not posited beside it. The burden is therefore not to invent a new electromagnetic U(1) and then identify it with the derived one; it is to explain why the U(1) already forced by transport would *not* be the one electromagnetism uses — and the structure-group audit of §3 is exactly the reading-off that settles it.

Fourth — a second circle is surplus structure unless it works. A decoupled electromagnetic U(1) would be a second independent phase circle in the candidate-layer sector. By no-pre-individuation, such a circle is not admissible merely because it is mathematically possible: it must perform independent physical work — it must carry some admissible distinction unavailable to the substrate phase. This converts the refuter into its exact form: *exhibit an admissible electromagnetic phase invariant that cannot be represented as substrate transport holonomy*. Without such an invariant, the second-circle hypothesis is not the conservative alternative it appears to be; it is extra structure of exactly the kind the programme forbids. The burden does not sit symmetrically between GI and its negation — the negation is the one that must produce something.

What the audit does and does not do. It does not prove GI, and the marker does not change: [Conditional], as carried. What it establishes is that the condition is auditably narrow. GI is not a casual identification of two similarly named objects; it is the claim that the one compact path-holonomy structure forced by the substrate is the same path-holonomy structure through which electromagnetic charge is defined. The disjunction left standing is exact: either electromagnetism uses the derived transport circle, or it requires a second physically independent circle whose admissible invariant content must be exhibited. The asymmetry of burden — the alternative must name its invariant — is the same burden-shift posture the U(1) paper adopts for its closure argument, and it is the audit's product. [Motivation plus burden-shift; GI's marker is unchanged by it.]

4. What Carries the Representation, and Why It Is a Representation

The conventional presentation begins "let the field transform as $\psi \rightarrow e^{iq\theta}\psi$ " — assuming both the carrier and the form of its response. The programme can ground both.

The carrier. Admissible transport carries the candidate-layer structure along admissible paths [Inherited, U(1) paper §4]. The object that responds to phase transport is that candidate-layer structure: the structure whose interference the Born chain quantifies, the structure on which the kernel $W(P, P') = e^{i(\theta(P) - \theta(P'))}$ acts. Call a species of such structure a *carrier*. The question of this paper is how a carrier may consistently respond to the transport phase.

Response supervenience on holonomy. Before ρ can be written as a function of θ at all, a premise must be made explicit that conventional treatments cross silently and this paper's own grounding forbids crossing silently. The substrate phase is path-dependent — that is the structural form §3 leans on — and two distinct admissible paths can accumulate the same phase θ . Writing the response as $\rho(\theta)$ presupposes that the carrier's response depends only on the

accumulated holonomy value, not on which path realizes it: *paths of equal holonomy produce equal carrier response*. This is a distinct claim, and it is forced by the same instrument that closes the paper's other gaps. Suppose two admissible paths of equal holonomy produced different carrier responses. The difference would be a physical distinction between the two transports; yet every admissible invariant the phase structure carries — the holonomy, and everything composition builds from holonomies — agrees between them by hypothesis. A response difference between invariant-identical transports is a distinction located where no admissible invariant changes: an admissibly-unwitnessed distinction, forbidden by the Conservation Lemma. The response therefore factors through the holonomy, and $\rho: U(1) \rightarrow U(1)$ is well-defined. [Proven, conditional on the Conservation Lemma [Inherited]. The premise is named *response supervenience on holonomy*; §12 verifies that it is catalogue-level, since it is the one new input whose grain could threaten the independence claim.]

Why a homomorphism. Let $\rho(\theta)$ denote the action on a carrier of transport accumulating phase θ . Two constraints are inherited, not chosen. First, *composition*: transports compose, and the action of a composite transport is the composite of the actions —

$$\rho(\theta_1 + \theta_2) = \rho(\theta_1) \rho(\theta_2),$$

since performing two admissible transports in sequence is itself an admissible transport whose accumulated phase is the sum [Inherited; the composition law of the U(1) paper §5]. Second, *reversibility*: every admissible transport has an admissible inverse and acts on amplitudes by unit-modulus factors [Inherited]. Invertibility and norm-preservation are distinct properties in general — reversible is not unitary without more — but for the one-dimensional carriers at issue, unit-modulus action *is* norm-preservation, so the gap closes for free in exactly the case the argument needs. Each $\rho(\theta)$ is therefore invertible and norm-preserving — unitary —

$$\rho(\theta)^{-1} = \rho(-\theta), |\rho(\theta)v| = |v|.$$

A norm-preserving, composition-respecting response is exactly a unitary representation of the phase group. The representation property is therefore not a modelling choice; it is the transcription of inherited transport structure into the carrier's response. [Proven, given the inherited inputs.]

One-dimensional carriers first. The elementary case — and the one electric charge inhabits — is the one-dimensional carrier: a single complex amplitude responding to phase. For a one-dimensional unitary representation, $\rho(\theta)$ is a unit-modulus complex number for each θ , and the composition law forces the exponential form

$$\rho(\theta) = e^{i q \theta}$$

for some real parameter q — *provided ρ is continuous*. That proviso is a genuine gap, not a pedantic one, and the next section closes it with the programme's own tools rather than by appeal to physical reasonableness. Higher-dimensional carriers decompose, once continuity is in hand, into one-dimensional characters (U(1) is abelian; its finite-dimensional continuous unitary representations are direct sums of characters, and the infinite-dimensional case follows by the

spectral theorem for continuous unitary representations), so the one-dimensional case is the general case for the abelian sector. [Proven, conditional on §5.]

5. Continuity of Response, Forced — Closing the Discontinuous-Character Loophole

The textbook classification — "the one-dimensional unitary representations of $U(1)$ are $\theta \mapsto e^{iq\theta}$, $q \in \mathbb{Z}$ " — is true only of *continuous* representations. Discontinuous characters of $U(1)$ exist: the circle group admits unmeasurable homomorphisms to itself, constructible from a Hamel basis of \mathbb{R} over \mathbb{Q} , which respect the composition law exactly while being discontinuous everywhere and integer-labelled nowhere. Conventional physics suppresses the loophole by declaring continuity physically obvious. This programme does not have that license — and does not need it, because it has a principle whose job is exactly this.

Continuity of the Carrier Response

Every admissible carrier response ρ is continuous on $U(1)$.

The argument runs in two layers, and the layering is itself load-bearing: the Conservation Lemma directly delivers a statement about the admissible grain, and topological continuity on all of \mathbb{T} is then *derived* from it rather than conflated with it.

Layer (a) — admissible continuity. Suppose the response had a gap at the admissible grain: phases θ admissibly inseparable from θ_0 — within the minimum distinguishable phase, which by the $U(1)$ paper's observability/catalogue distinction (its §2) bounds what any local comparison can resolve, with composition-amplification (its §7.4) the only admissible route below that scale, a route available only between path-realized holonomies — at which $\rho(\theta)$ differs from $\rho(\theta_0)$ by more than some fixed amount. Then the carrier's behavior changes across configurations between which no admissible invariant changes: a response gap at an invariant-free location. By the Conservation Lemma, physical structure may change only where some admissible invariant changes; the gap is an admissibly-unwitnessed distinction, and no-pre-individuation forbids the carrier from carrying it. So: no response gap without an invariant change — admissible continuity.

An objection must be met head-on rather than in passing, because it turns the $U(1)$ paper's own machinery against this argument. Composition-amplification exists precisely to make sub-grain phase differences witnessable: compose enough copies of a transport and the difference $n(\theta - \theta_0)$ climbs above the distinguishable scale. If amplification can in principle separate any two distinct phases, no two distinct phases are unconditionally inseparable, and layer (a) threatens to be vacuous — every would-be response gap would sit at a location some composed invariant could distinguish. The resolution is the one the $U(1)$ paper already litigated for its closure argument: the Conservation Lemma runs on *realized* witnesses, not on witnesses constructible in principle. An amplification that has not been admissibly performed is exactly the "unrealized limit" that the $U(1)$ paper's closure argument declines to count — its architecture line "no amplification witness at unrealized limits" is this distinction, applied there to the catalogue boundary and here to the

response. Between admissibly-inseparable phases whose separating amplification is unrealized, no admissible invariant changes; a response gap there remains unwitnessed, and the Lemma forbids it. The refuter of §13 inherits the sharpened form: name a *realized* witness.

[Proven, conditional on the Conservation Lemma [Inherited] and the U(1) paper's §2 and §7 results [Inherited]; the paper leans here on the U(1) paper's most delicate balance — finite distinguishability coexisting with a continuum catalogue — and the lean is on the exact mechanism, including its realized/in-principle witness distinction, cited, not on the slogan.]

Layer (b) — topological continuity. Admissible continuity is a statement at and below the admissible grain; topological continuity of ρ as a character of \mathbb{T} is a statement at all scales, and the bridge must be built rather than gestured at. It is built from two facts. First, the homomorphism property globalizes local behavior: for a character, $\rho(\theta + \varepsilon) / \rho(\theta) = \rho(\varepsilon)$, so the behavior of ρ near the identity controls its behavior near every point, and continuity at one point is continuity everywhere; gaplessness near the identity is exactly what layer (a) supplies. The step from gaplessness to continuity is itself classical and citable: a homomorphism into \mathbb{T} that maps some neighbourhood of the identity into a sufficiently small arc is continuous, by the divisibility argument — $\rho(\varepsilon) = \rho(\varepsilon/2)^2$, with the branch of each successive square root pinned inductively by the small-arc hypothesis — and the small-arc hypothesis is precisely what layer (a) supplies. Second, a mathematical backstop converts the diagnosis into something theorem-shaped: every *measurable* character of U(1) is continuous — a classical result of Weil–Banach type — and the discontinuous characters are precisely the non-measurable ones, constructed via a Hamel basis of \mathbb{R} over \mathbb{Q} , that is, via the axiom of choice, and specifiable by no explicit construction whatsoever. A response that no admissible construction can specify is not a candidate physical response; the unwitnessed-distinction diagnosis of layer (a) is, in mathematical dress, exactly the statement that the discontinuous characters live outside everything constructible. The two layers together force: every admissible carrier response is a continuous character of \mathbb{T} . [Proven, given (a) and the homomorphism property of §4.]

The Lemma's application here has a different grain than at source, and the difference is worth stating once. At source the Lemma removed a boundary in the *catalogue* — a difference in admissibility between a limit point and its neighbours. Here it removes a jump in the *response* — a difference in carrier behavior between admissibly-inseparable phases. The common form is exact: in both cases a physical difference is posited at a location where the distinguishability data carry no difference, and in both cases the difference is therefore not admissible structure. The discontinuous characters of U(1) are thereby in the same position as the dense proper subgroups dispatched at source: legitimate as mathematics — indeed constructible only by choice, hence not constructible at all in the operative sense — and inadmissible as physics, for the same reason: they encode distinctions the admissible structure cannot see. [Proven, conditional on the Conservation Lemma and the U(1) paper's closure and amplification results, all [Inherited].]

With continuity forced, the classification of §4 is complete: every one-dimensional admissible carrier response has the form $\rho(\theta) = e^{iq\theta}$ with q real and ρ continuous. What remains is the constraint that compactness places on q .

6. The Compactness Constraint — Integer Labels

On the compact circle, the phases θ and $\theta + 2\pi$ are the same point of $U(1)$ — not nearby points, not conventionally identified points, but one point, by the exact periodicity of §2. A carrier response is a function on $U(1)$; it cannot assign two values to one point. Therefore

$$\rho(\theta + 2\pi) = \rho(\theta) \text{ for all } \theta.$$

Substituting the form forced by §§4–5:

$$e^{iq(\theta + 2\pi)} = e^{iq\theta},$$

and cancelling the common factor,

$$e^{2\pi iq} = 1.$$

The unit-modulus exponential equals 1 exactly when its argument is an integer multiple of 2π :

$$q \in \mathbb{Z}, q = 0, \pm 1, \pm 2, \pm 3, \dots$$

The Charge Lattice

The admissible one-dimensional carrier responses to compact $U(1)$ phase are exactly $\rho_q(\theta) = e^{iq\theta}$ with $q \in \mathbb{Z}$. The response labels form the integer lattice.

[Proven, conditional on GI (§3) and the inherited inputs of §§2, 4, 5.]

The mechanism deserves a sentence of plain statement, because it is the entire physical content. On the non-compact line, the consistency condition $\rho(\theta_1 + \theta_2) = \rho(\theta_1)\rho(\theta_2)$ constrains the *form* of the response but not the label: every real q is consistent, because no two distinct phases are ever identified. On the circle, one global identification — $\theta \equiv \theta + 2\pi$ — is added, and that single closure condition discretizes the entire label spectrum. Compactness converts a continuum of consistent couplings into a lattice. The $U(1)$ paper derived the circle; the circle does the rest.

7. Charge as the Representation Label

What physical quantity is the integer q ? Under GI, the answer is fixed by what electric charge *is* in gauge theory: charge is the coupling of a field to the gauge phase — the rate at which the field's own phase turns per unit turn of the gauge phase. The transformation law

$$\psi \rightarrow e^{iq\theta} \psi$$

is not an equation charge appears in; it is the definition of the charge of ψ . The representation label of §6 and the electric charge are the same object under the identification GI makes: the phase the carrier responds to is the compact transport phase, so the carrier's coupling to it is its charge, and the coupling is integer-valued.

Compact $U(1)$ (derived) + GI (named premise) $\Rightarrow q \in \mathbb{Z} \Rightarrow$ charge is lattice-valued.

The geometric reading is winding. A carrier of charge q winds q times in its own phase per single circuit of the transport phase. Integer winding is the only consistent meshing of two circles: a carrier winding 1.37 times per circuit does not return to its own starting configuration when the transport phase closes its loop, and the mismatch is a genuine inconsistency — the carrier would carry two values at one point of $U(1)$, exactly what §6 forbids. Charge, on this account, is not a primitive attribute attached to matter. It is the integer winding number of the carrier's admissible response to compact phase: a topological label, which is *why* it is exact, conserved, and identical across every particle of a species — integers do not drift, and there is nothing between the rungs for a charge to drift to. One caveat keeps the reading at its correct strength: exactness-by-topology holds given GI and given that the carrier remains a continuous character under all admissible dynamics, so the conservation claim inherits exactly the conditionality of §§3 and 5, not a stronger one. The extraordinary measured equality of electron and proton charge magnitudes — agreement to a part in 10^{21} — is, under this reading, not a fine-tuning but the equality of the integers 1 and 1. [Proven, conditional on GI; the topological reading is interpretation, carried at the same status.]

8. Why Non-Compact Phase Fails — and the Inference Run Backwards

The counterfactual sharpens the result. Suppose the phase to which charge couples were valued in \mathbb{R} — accumulating without periodic identification, the line rather than the circle. Then no closure condition exists: θ and $\theta + 2\pi$ are distinct phases, $\rho(\theta + 2\pi) = \rho(\theta)$ is not required, and the consistency analysis of §6 terminates at $\rho(\theta) = e^{iq\theta}$ with q an arbitrary real. Every coupling is admissible. Charge is continuous; carriers with charges 1, 1.0000003, and π coexist consistently; the lattice is gone.

This failure can be run in both directions, and the backward direction is independent evidence.

Forward (this paper): derived compact $U(1) + \text{GI} \Rightarrow$ lattice charge — matching observation.

Backward (observational): observed charge is lattice-valued to extreme precision, and the lattice is exactly what non-compact phase cannot produce. The empirical discreteness of charge is therefore observational support for the compactness of the phase charge couples to — and, through GI, for the $U(1)$ paper's compact-phase result. The Born-rule and interference phenomenology constrain the transport phase; charge discreteness constrains the gauge phase; under GI these are one structure receiving support from two independent empirical directions.

Two cautions keep the backward inference honest, in the spirit of the $U(1)$ paper's empirical-standing discipline. First, the support is shared: conventional physics also assumes compact $U(1)$, so observation favours compactness over \mathbb{R} without discriminating *derived* compactness from *assumed* compactness — the derivation must stand on its own argument. Second, charge discreteness discriminates the circle from the line but not from the line's other rivals: a dense proper subgroup of the circle (the rational-angle catalogue of the $U(1)$ paper's §7.5) would also yield discrete-looking charge phenomenology at any finite resolution, so the case against dense gapped phase remains structural, not empirical — exactly as at source. The observational lever

bears wholly against \mathbb{R} ; the U(1) paper's closure argument carries the rest. [Observational consistency; no confirmatory weight claimed beyond the stated direction.]

9. Fractional Charges — Embedding, Not Exception

Quarks carry charges $\pm 1/3$ and $\pm 2/3$ of the electron charge. Does the integer lattice of §6 fail against the most familiar facts of hadronic physics?

It does not, and the resolution is standard, but its status under this derivation should be stated exactly. The lattice result says the admissible charges are integer multiples of *some* elementary unit q_0 — the generator of the lattice. It does not say which observed particle carries q_0 . If the true elementary unit is one-third of the electron charge, then quarks occupy rungs 1 and 2 of the lattice (in down-type and up-type magnitudes) and the electron occupies rung 3; every observed charge is an exact integer multiple of $q_0 = e/3$, and the lattice is intact. What requires explanation is then not fractionality — there is none, measured against the true unit — but the *occupancy pattern*: why leptons sit at multiples of three while quarks occupy the intermediate rungs, and why the combination rules (three quarks to a baryon) restore multiples of three.

In the Standard Model that pattern is the business of the embedding: the electromagnetic U(1) sits inside $SU(3) \times SU(2) \times U(1)$ via

$$Q = T_3 + Y/2,$$

and the observed charges follow from the representation assignments of the fermions under the full group — with anomaly cancellation constraining those assignments tightly (§11). Deriving the embedding, and hence the occupancy pattern, is the multi-Fold programme's burden: it requires the non-abelian sectors, which the U(1) paper explicitly marks out of scope. The present paper's claim is therefore exactly bounded: *discreteness is derived here; the unit and the occupied rungs are downstream*. The lattice does license one crisp standalone prediction, and it is stated as one: no charge will ever be observed that is not an integer multiple of $e/3$. Fractional charges are not a counterexample to the lattice — they are evidence about where its generator sits. [Proven that the lattice accommodates the observed spectrum under $q_0 = e/3$; the derivation of the embedding is [Open] at programme level, assigned to the multi-Fold arc.]

10. The Unit of Charge — What This Paper Does Not Fix

A lattice has two parameters: its structure and its spacing. This paper derives the structure — integer multiples of a generator — and fixes nothing about the spacing. The physical magnitude of the elementary charge, equivalently the electromagnetic coupling strength e (and with it the fine-structure constant $\alpha = e^2/4\pi\epsilon_0\hbar c$), is invisible to every argument made here: the representation analysis of §§4–6 is unchanged under any rescaling of the unit, because integrality is a statement about ratios of charges, not about their absolute size.

This is a feature of the derivation's grain, not an omission to be papered over. Topological arguments fix discrete labels; they do not fix continuous couplings — the winding number of §7

says how many times the carrier turns, not how strongly the turning costs. Within the programme, the magnitude of the coupling is the business of the fine-structure arc, which addresses α by a separate route; the present result and that one are complementary and independent — the lattice from transport-and-closure structure, the spacing from the coupling derivation — and neither leans on the other. A reader should therefore not look to this paper for the value of e , and should not count its absence against the lattice result: the question "why is charge discrete?" and the question "why is the unit this size?" are different questions, and this paper answers only the first. [The division is definitional; the α derivation is carried at its source's status.]

11. Comparison with the Conventional Routes

The claim of novelty must be calibrated exactly, because the final inferential step here is textbook and pretending otherwise would be the wrong kind of originality.

What is standard. That compact $U(1)$ gauge group \Rightarrow integer charge lattice is classical representation theory, stated in any gauge theory course. The present paper adds nothing to that implication; §§4–6 execute it with the programme's own grounding (the carrier identified, the homomorphism property derived from transport structure, the continuity forced rather than assumed), but the implication itself is old.

What is supplied. What has never been available is the premise. Conventionally, the compactness of electromagnetic $U(1)$ is an assumption — and the recognized cost of leaving it one is exactly why the alternative routes exist:

Grand unification derives compactness by embedding $U(1)$ in a compact simple group ($SU(5)$, $SO(10)$); the embedding discretizes charge through the parent group's representation theory. Price: the parent group's new gauge bosons, its breaking sector, and proton decay — unobserved at the predicted scales.

Dirac monopoles derive the quantization condition $qg = 2\pi n$ from single-valuedness of the wavefunction around a monopole; one monopole anywhere quantizes all charge everywhere. Price: the monopole — unobserved.

Anomaly cancellation constrains the hypercharges of the Standard Model fermions through the requirement that gauge anomalies vanish; with the observed fermion content, the constraints are strong enough to fix charge ratios up to discrete choices. This is the strongest conventional argument and the one a fair comparison must foreground. Its character differs from the present route: it is a consistency condition *within* an assumed gauge structure and fermion content, fixing ratios among the charges present, rather than an explanation of why charge is lattice-valued prior to any particle content. It also leans on the assumed compactness of the ambient groups. The two routes are therefore complementary rather than competing: the present derivation supplies lattice-valuedness from below; anomaly cancellation disciplines the occupancy of the lattice by the actual fermion spectrum — which is precisely the §9 division between discreteness (here) and occupancy pattern (downstream).

The present route. Compactness is derived — from reversibility, unbounded composition, and no-pre-individuation, in the U(1) paper — and charge quantization then follows with no unobserved structures introduced. The honest statement of cost: this route's price is not new particles but new premises — the substrate architecture of the programme, carried at the corpus's statuses, and the gauge identification GI of §3, carried as [Conditional] with a named refuter. Whether that price is lower than a monopole is for the reader; that it is a *different* price, paid in foundations rather than in unobserved physics, is the claim. The contrast, in two lines:

Standard route: assume compact U(1) \Rightarrow charge lattice. VERSF route: derive compact U(1) + GI \Rightarrow charge lattice.

The mathematical implication is old; the proposed source of compactness is new. The paper's success should therefore be judged at the right point: not on whether compact U(1) implies integer representation labels, which is standard, but on whether the U(1) paper's phase result genuinely earns the compactness, and whether the gauge identification audit (§3.2) succeeds in identifying that compact phase with electromagnetism. [Calibration; no marker.]

12. Position in the Programme — Independence from the Assignment Node

The U(1) paper's architecture diagram carries one arrow unearned: the assignment question — whether the holonomy assignment $L \mapsto h(L)$, and not only its value space, is D-determined — carried at [Open], blocking the generation chain (support supervenience, dyadic loading). The Assignment Paper addresses that node, concluding D-determination up to gauge conditionally on its named Transport-Completeness premise; the node stands at [Conditional] at that paper's status. The present result consumes the node under neither status, and the independence is worth establishing explicitly rather than leaving to inspection. Every input consumed above is catalogue-level: the compactness, connectedness, and exact periodicity of the value space (§2); the Conservation Lemma (§5); the composition and reversibility of transport (§4). At no point does any argument require that a *particular* path carry a *particular* holonomy — the representation analysis quantifies over the phase group, not over the assignment, and the closure condition of §6 is a condition on functions on the value space. Two holonomy assignments sharing the catalogue U(1) but differing on paths would yield identical charge lattices: the lattice is a property of the circle, and both assignments fill the same circle.

One input requires the audit to be re-run visibly rather than assumed to survive. §4 carries response supervenience on holonomy as a named premise: equal-holonomy paths produce equal response. Is it catalogue-level? It is, and the verification is short. The premise quantifies over *pairs of paths with equal holonomy* — it constrains the response's dependence structure (which invariant the response factors through) without ever consuming *which* path carries *which* holonomy. Holonomy equality is a relation internal to any given assignment: two rival assignments sharing the catalogue would each induce their own equal-holonomy pairs, and the premise binds the carrier identically under either. It therefore constrains carriers, not assignments, and the independence claim survives the repair.

The corrected programme diagram, with this paper's branch attached at its true point of dependence:

```

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

H infinite + subgroup ⇒ H dense in  $\mathbb{T}$ 

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

Dense + closed ⇒ admissible phase structure =  $\mathbb{T}$ 
Reversible ( $\mathbb{C}$ -realized), single-Fold ⇒ exactly  $U(1)$ 
├── U(1) ⇒ Born obligation discharged
├── U(1) + GI [Conditional] ⇒ charge lattice ← this paper
└── U(1) + assignment  $L \mapsto h(L)$  D-determined up to gauge
    [Conditional – Assignment Paper, on Transport-Completeness of
D]
    ⇒ support supervenience ⇒ dyadic loading

```

The consequence for the programme's sequencing is the quiet significance of this paper: the charge lattice is the first piece of Standard-Model-facing physics in the corpus that does not wait on the assignment node. The generation chain waits on Transport-Completeness specifically; the charge chain does not. A failure of Transport-Completeness in the Assignment Paper would leave the present result standing untouched; conversely, nothing here advances the assignment question, and no claim is made that it does. [Proven — the independence — given the catalogue-level character of every consumed input.]

12.1 Why the Charge Result and the Assignment Result Are Complementary

The Assignment Paper and the present paper consume different aspects of the phase programme, and the difference is exact rather than approximate.

The Assignment Paper concerns the map from admissible histories to holonomies: whether the gauge-invariant phase assignment carries information beyond admissible distinguishability. Its machinery is the Holonomy Dilemma — a difference in closed-loop holonomy is either admissibly witnessed, hence recorded in D , or unwitnessable, hence not physical structure — run under its named Transport-Completeness premise. The present paper consumes none of that machinery. Its representation-theoretic argument depends only on the compactness of the phase catalogue, as §12 establishes input by input, and would survive unchanged if Transport-Completeness fell.

The relationship is nevertheless informative, because the two papers meet at the Gauge Identification premise from opposite directions. If the Assignment Paper stands, the transport phase is not merely compact but physically distinguished: its gauge-invariant content is already fixed by admissible comparison structure (§3.1). The charge paper then asks whether

electromagnetism uses that already-distinguished phase. One paper reduces the freedom available to rival phase structures; the other shows what follows if electromagnetism adopts the phase structure that remains.

Taken together — and conditionally on each paper's own status — the two transform the interpretation of GI. It ceases to be an identification between two similarly named $U(1)$ groups and becomes the hypothesis that electromagnetism uses the unique compact, D -determined phase structure the substrate already forces. The hypothesis remains [Conditional]; what the conjunction changes is what its negation would cost. A rival must then posit a second physically effective phase structure beside a derived one that already does the work — and must exhibit the admissible invariant content that justifies it. [Complementarity; each input carried at its source status.]

13. What Would Refute This

Each load-bearing element has a precise refuter; naming them makes the result a defined target.

Against GI, the gauge identification (§3, §3.2) — the load-bearing refuter. Exhibit an admissible electromagnetic phase invariant that cannot be represented as substrate transport holonomy. By the audit of §3.2 this is the exact and only form a refutation can take: a second independent circle is admissible structure only if it performs independent physical work, so the challenger must name the invariant the transport phase cannot carry — equivalently, demonstrate within the Maxwell-sector papers that the structure group of the Maxwell connection constructed there is not the derived transport T . Standing: no such invariant is known; electromagnetic phase phenomenology — interference, Aharonov–Bohm holonomies, field strength as curvature — is holonomy phenomenology in form, and the single-Fold architecture supplies no second circle. The burden is the challenger's to name the invariant. The consistency audit against the Maxwell sector remains the named successor obligation; GI is the one [Conditional] this paper adds to the corpus, and the one place to press.

A stronger form of the refuter becomes available conditionally — if the Assignment Paper stands on its Transport-Completeness premise. What the strengthening adds is the closure of an inversion. Without the Assignment Paper, a challenger has an escape route: concede the catalogue but argue that the *transport* phase is the idle auxiliary and the *electromagnetic* phase the physically operative one — inverting which circle owes the justification, so that GI's burden-shift points the wrong way. With D -determination, that inversion closes: the transport phase is established as physically distinguished — its gauge-invariant content fixed by admissible comparison itself — so a rival circle cannot present itself as the real structure displacing a free one. It can only be a *second* effective structure, and the demand sharpens accordingly: a successful refuter must exhibit admissible invariant content carried by the electromagnetic phase that is not already carried by the D -determined transport phase. Merely renaming the phase, or assigning it a different mathematical origin, does not suffice; the invariant must generate distinct observable structure — a difference in provenance with no admissible witness is no difference at all, by the same standard the rest of the programme is held to. Without such content, the second phase contributes no admissible distinction beyond transport structure and falls under the same no-pre-individuation pressure that removed unsupported structure elsewhere in the programme.

The unconditional refuter above remains the operative one; this strengthened form is recorded at the Assignment Paper's status (§3.1, §12.1).

Against the continuity forcing (§5). Exhibit a *realized* admissible witness — an admissible invariant, actually performed rather than constructible in principle, that changes across the would-be discontinuity of a carrier response. Such a witness would make the discontinuous response admissible structure, the discontinuous characters would re-enter, and the integer classification would fail. An in-principle amplification does not suffice: the Conservation Lemma runs on realized witnesses, exactly as the U(1) paper's closure argument declines to count unrealized limits (§5). Standing: the refuter has the same shape as the U(1) paper's closure refuter, and the same burden — name the realized witness.

Against response supervenience (§4). Exhibit two admissible paths of equal holonomy producing distinguishably different carrier responses — a response difference witnessed by some admissible invariant that the holonomy does not carry. Such a pair would show the response not factoring through the phase; ρ would not be a function on U(1), and the classification would not begin. Standing: any witnessing invariant would be admissible structure beyond the holonomy in the candidate-layer sector, which the single-Fold architecture does not supply; the burden is the challenger's to name it.

Against the representation premise (§4). Exhibit an admissible carrier whose response to composed transport is not the composition of its responses — breaking the homomorphism property — or one whose response is not norm-preserving despite reversible transport. Standing: both properties are transcriptions of inherited transport structure; a counterexample would contradict the inherited inputs at source.

Against the conclusion, empirically. Observe a charge off the lattice: any two charges whose ratio is irrational, or a continuously tunable charge, refutes the result outright — and refutes it upstream, since under GI an off-lattice charge implies non-compact or gapped phase, contradicting the U(1) paper. Standing: Millikan-type searches for fractional (non-thirds) charge, and the part-in- 10^{21} electron–proton balance, currently sit entirely on the lattice. This is the live falsification channel, and it bears on the whole chain at once.

Inherited exposure. The result is additionally conditional on every input of the U(1) paper at its source status — no-pre-individuation, reversibility, finite-holonomy-impossibility, single-Fold binarity — and falls with any of them. The exposure is shared with the Born chain and is not re-audited here.

14. What the Paper Establishes

Established (conditionally, with the conditions named):

- The gauge identification audit (§3.2): four structural coincidences — shared holonomy form (with field strength as curvature, so force phenomenology is holonomy phenomenology), shared carrier, the Maxwell sector's transport character, and the no-pre-individuation burden against an idle second circle — narrowing GI's conditional to an

exact disjunction and yielding the refuter's invariant form. [Motivation plus burden-shift; GI's marker remains [Conditional].]

- Response supervenience on holonomy: equal-holonomy paths produce equal carrier response, forced by the Conservation Lemma — the well-definedness premise conventional treatments cross silently, here named, argued, and marked (§4). [Proven, conditional on the Lemma [Inherited].]
- The carrier and the homomorphism property of its response, derived from inherited transport structure rather than assumed; unitarity from unit-modulus reversibility, with the invertibility/norm-preservation gap closed explicitly for the one-dimensional case (§4). [Proven, given inherited composition and reversibility.]
- Continuity of every admissible carrier response, in two layers — admissible continuity forced directly by the Conservation Lemma; topological continuity derived from it via the homomorphism property, with the Weil–Banach measurability backstop converting the unwitnessed-distinction diagnosis into theorem-adjacent form: the discontinuous characters are exactly the non-measurable, choice-constructed ones, specifiable by no admissible construction (§5). [Proven, conditional on the Lemma and the U(1) paper's closure and amplification results, all [Inherited].]
- The charge lattice: admissible one-dimensional responses to compact U(1) phase are exactly $\rho_q(\theta) = e^{iq\theta}$, $q \in \mathbb{Z}$ (§6); under GI, q is electric charge, and charge is lattice-valued (§7). [Proven, conditional on GI.]
- The topological reading: charge as integer winding of the carrier response, whence its exactness, conservation, and species-wise identity (§7). [Interpretation, at the same status.]
- The backward inference: observed charge discreteness as observational support for compact phase, with the stated cautions — support shared with assumed-compactness physics, and silent against dense proper subgroups (§8). [Observational consistency only.]
- The exact bounding of the fractional-charge question: the lattice accommodates the observed spectrum with generator $q_0 = e/3$; occupancy pattern and embedding are downstream, assigned to the multi-Fold arc (§9).
- The lattice/spacing division: structure derived here, unit not fixed here, by the grain of topological argument; the unit is the fine-structure arc's business (§10).
- The calibrated novelty claim: the representation step is textbook; the contribution is the derived premise — compactness from below — replacing the unobserved structures that conventionally supply it; anomaly cancellation acknowledged as complementary, disciplining lattice occupancy rather than lattice-valuedness (§11).
- Independence from the assignment node: every consumed input is catalogue-level — including, by explicit re-audit, the response-supervenience premise, which conditions only on holonomy equality and never on which path carries which holonomy — so the charge chain does not wait on the generation chain's remaining question (§12). [Proven.]

Not established (open or out of scope):

- the gauge identification GI itself (§3) — the paper's one new [Conditional]; its verification against the Maxwell-sector papers is the named successor obligation, and its refuter is named (§13);

- the value of the elementary charge and the coupling strength (§10) — the fine-structure arc's burden;
- the specific charge assignments of the Standard Model fermions — the occupancy pattern, the embedding $Q = T_3 + Y/2$, and the role of anomaly constraints (§9, §11) — the multi-Fold arc's burden;
- the assignment question and everything downstream of it (§12, §12.1) — addressed by the Assignment Paper at [Conditional] on its Transport-Completeness premise, with the holonomy principle [Imported-External]; unconsumed and unadvanced here;
- all inherited inputs, at their source statuses (§13).

The honest summary: this paper executes a textbook implication whose premise has never before been available as a derived result, and adds exactly one new conditional to the corpus in doing so. Compact phase was derived in the U(1) paper; given the gauge identification, the compactness discretizes every consistent coupling, and electric charge — the coupling — is forced onto the integer lattice. The lattice is derived; the unit, the occupancy, and the identification's audit are named downstream work.

15. Conclusion

The Standard Model contains the charge lattice as an input. Grand unification buys an explanation with unobserved bosons; the monopole argument buys one with an unobserved particle; anomaly cancellation disciplines the lattice's occupancy but presupposes the compact ambient structure. The cheapest explanation has always been hiding in the representation theory: if the phase that charge couples to is a compact circle, charge is integer-valued, by an argument of a few lines. The reason that argument has never settled the question is that the compactness was always an assumption — the very thing in need of explanation, relocated.

The U(1) paper removed that status. Compactness is, within this programme, a theorem: reversible transport acts by unit-modulus factors, the catalogue can carry no unwitnessed edge, and the phase structure is forced to the full circle U(1) — derived, conditionally, from finite distinguishability and the discipline that admissible structure is exhausted by admissible distinguishability. The present paper carries that theorem the last few steps: the carrier's response is a representation because transport composes and reverses; the response is continuous because a response gap at an invariant-free location is an unwitnessed distinction the Conservation Lemma forbids; and on the circle, one global identification — $\theta \equiv \theta + 2\pi$ — discretizes every consistent coupling. The hand of the dial may drive its gear once, twice, an integer number of times per revolution; it cannot drive it 1.37 times and close the loop. Charge is the gear ratio, and the ratios are integers.

One conditional separates this from a closed result, and it is named rather than crossed: that the phase of electromagnetism is the phase of transport — one circle, not two. The structural case is strong (the substrate phase is born path-dependent, in the connection form; the single-Fold architecture supplies exactly one phase), and the audit against the Maxwell sector is the defined next step. Granting it, the oldest discreteness in physics stops being a brute fact: charge is quantized because phase is compact, and phase is compact because a finite world, composing

without bound, can sustain no structure that nothing in it could ever witness. The lattice was never an extra law. It is what the circle leaves room for.