Fermion Structure and Particle Classification from Codec Geometry

Juha Meskanen

June 2026

Abstract

We prove three exact theorems about the density matrix of a fermion in equal superposition across n sites, within the codec framework of Papers I–VII. Let $F = \operatorname{diag}(\rho)$ be the fermionic diagonal and $B = \rho - \operatorname{diag}(\rho)$ the bosonic off-diagonal. Then:

$$\|F\| = \frac{1}{\sqrt{n}} \quad \text{(Universal Fermion Norm)},$$
$$\|B\| = \sqrt{\frac{n-1}{n}} \quad \text{(Universal Boson Norm)},$$
$$\|F\|^2 + \|B\|^2 = 1 \quad \text{(Conservation Law)}.$$

All three results are phase-independent and hold for arbitrary $n \geq 1$. The winding number m of the phase pattern determines the internal symmetry of the fermion's compression residual.

The central claim of this paper differs from the Standard Model in a foundational respect: bosons are not fundamental particles. They are the codec's internal record of fermion transitions, and are never directly observed. Every physical observation is a fermion observation. The particle classification developed here is therefore a *fermion* classification: the pair (n, m) describes the fermion's complexity class, and what Standard Model language calls a "gauge boson" is the name we give to correlations between fermion events that would otherwise require an unexplained interaction term. Under this reframing, the open problems of the framework reduce to purely fermionic questions: what distinguishes the charged lepton generations, why fermions carry the charges they carry, and why their trajectories produce the force laws they produce. A combinatorial argument for a 4-bit codec unit per fermion hop is noted but is explicitly not connected by any derivation to the observed lepton mass gaps, which are unequal (7.69 and 4.07 bits) and do not both round to a common unit.

Contents

1 Introduction

The codec framework of Papers I–VII treats the universe as a finite static bitstring of $n \approx 2^1 84$ bits. Fermions are localised excitations — pixels — encoded as normalised complex wavefunctions. Bosons are compression residuals: the off-diagonal entries of the density matrix $\rho = |\psi\rangle\langle\psi|$ that remain after the per-site fermionic diagonal is removed.

Paper VII showed that a two-site fermion hop produces a compression residual B with $\|B\| = \sin(2\theta)/\sqrt{2}$, purely antisymmetric structure, and a universal $-\pi/2$ phase encoding the Born rule exactly. That result is about the *fermion's motion* — the boson is what the codec records when a fermion moves, not a separate particle that mediates the motion.

The present paper asks: what is the universal structure of the fermionic content itself, and how does it organise into the particle spectrum we observe?

The answer has two parts. First, three exact theorems establish the norm structure of any n-site fermion, culminating in a conservation law that subsumes all previous results. Second, the pair (n, m) — sites and winding number — classifies the fermion complexity classes. The classification is purely fermionic: no interaction terms, coupling constants, or gauge symmetries are introduced. What Standard Model language calls bosons are derived labels for correlations between fermion events, not independent objects.

A note on scope: gravity was derived in Papers IV–V at the geometric layer, independently of the wavefunction codec. That derivation is not revisited here. The relationship between the geometric and codec layers is addressed in the Discussion.

2 Setup

We follow the conventions of Paper VII exactly. A *fermion* is a localised excitation encoded as a normalised complex wavefunction $\psi \in \mathbb{C}^n$. The *density matrix* is $\rho = |\psi\rangle\langle\psi|$, with components $\rho_{ij} = \psi_i \psi_j^*$. We split ρ into two orthogonal parts:

$$F = \mathrm{diag}(\rho), \qquad B = \rho - \mathrm{diag}(\rho),$$

where F is the fermionic diagonal (per-site occupancy, real and non-negative) and B is the off-diagonal matrix encoding inter-site correlations.

An *n-site equal superposition* is

$$\psi_k = \frac{1}{\sqrt{n}} e^{i\phi_k}, \quad k = 0, \dots, n-1,$$

for arbitrary real phases ϕ_k. The natural phase patterns on an n-site ring are the discrete Fourier modes

$$\phi_k^{(m)} = \frac{2\pi m k}{n}, \quad m = 0, 1, \dots, \lfloor n/2 \rfloor,$$

where m is the *winding number* — the number of full phase rotations around the ring. These are the minimum-C_s states on the ring, ordered by ascending spectral complexity.

3 Three Exact Theorems

Theorem 1 (Universal Boson Norm)**.** For any n-site equal superposition with arbitrary phases,

$$\|B\| = \sqrt{\frac{n-1}{n}}.$$

Proof. $\rho_{ij} = (1/n)\, e^{i(\phi_i - \phi_j)}$, so $|B_{ij}|^2 = 1/n^2$ for all $i \neq j$. There are $n(n-1)$ off-diagonal pairs, giving $\|B\|^2 = n(n-1)/n^2 = (n-1)/n$. □ □

Theorem 2 (Universal Fermion Norm)**.** For any n-site equal superposition with arbitrary phases,

$$\|F\| = \frac{1}{\sqrt{n}}.$$

Proof. $F_k = |\psi_k|^2 = 1/n$ for all k, so $\|F\|^2 = \sum_k 1/n^2 = 1/n$. □ □

Theorem 3 (Fermion-Boson Conservation Law)**.** For any pure state ψ,

$$\|F\|^2 + \|B\|^2 = 1.$$

Proof. For any pure state, $\rho^2 = \rho$, so $\|\rho\|^2 = \mathrm{Tr}(\rho^2) = \mathrm{Tr}(\rho) = 1$. Since F and B occupy orthogonal subspaces, $\|\rho\|^2 = \|F\|^2 + \|B\|^2 = 1$. □ □

Remark (Connection to Paper VII). For $\psi(\theta) = \cos\theta\, e_i + i \sin\theta\, e_j$, $\|B(\theta)\|^2 = \sin^2(2\theta)/2$ (Paper VII, exact). Then $\|F(\theta)\|^2 = \cos^4\theta + \sin^4\theta$, and

$$\|F\|^2 + \|B\|^2 = \cos^4\theta + \sin^4\theta + 2\sin^2\theta\cos^2\theta = (\cos^2\theta + \sin^2\theta)^2 = 1.$$

Theorem 3 subsumes the Paper VII result as a special case.

Corollary 1 (Pauli exclusion)**.** For a localised fermion $\psi = e_j$: $\|F\| = 1$, $\|B\| = 0$. No compression residual is produced. Double occupancy has zero off-diagonal content and is informationally invisible.

Table 1 summarises the theorems for small n.

n	$\|F\| = 1/\sqrt{n}$	$\|B\| = \sqrt{(n-1)/n}$	$\|F\|^2$	$\|B\|^2$
1	1	0	1	0
2	$1/\sqrt{2}$	$1/\sqrt{2}$	1/2	1/2
3	$1/\sqrt{3}$	$\sqrt{2/3}$	1/3	2/3
4	1/2	$\sqrt{3}/2$	1/4	3/4

Table 1: Universal fermion and boson norms. All rows satisfy $\|F\|^2 + \|B\|^2 = 1$ exactly.

4 Fermion Spin Structure

The winding number m determines the symmetry of the fermion's compression residual. The half-integer spin of the fermion itself is encoded separately, in the universal phase structure established in Paper VII.

For any single fermion hop, the off-diagonal entries of B carry a universal $-\pi/2$ phase. Four successive hops accumulate a total phase of $4 \times (-\pi/2) = -2\pi$, returning the system to its original state. This is precisely the spinor double-cover structure of spin-1/2: a 2π rotation returns a fermion to its state only up to a sign, while a 4π rotation is the identity.

The conservation law therefore unifies two structures in a single codec object:

- The diagonal F carries half-integer spin, inherited from the $-\pi/2$ universal phase of the fermion hop.
- The off-diagonal B carries the internal symmetry class determined by the winding number m.

This is not a postulate; it follows from the imaginary encoding $\psi = f_0 + if_1$ and the definition of B, as shown in Paper VII.

5 Fermion Complexity Classes

The norms $\|F\|$ and $\|B\|$ depend only on n — the number of sites across which the fermion is delocalised. The internal structure of the compression residual is determined by the winding number m. Together, (n, m) defines a *fermion complexity class*.

The symmetry of the off-diagonal B under transposition determines whether the fermion can propagate freely or is confined:

- $B = -B^\top$ (antisymmetric, m odd for n=2): the fermion's residual has self-contained antisymmetric structure and can propagate as a free asymptotic state.
- $B = B^\top$ (symmetric, m=0): the residual has self-contained symmetric structure and propagates freely with scalar character.
- Mixed symmetry ($n \geq 3$, $m \geq 1$): the residual has no self-contained symmetry class. The fermion requires reference to its source configuration to be described — it cannot propagate as a free asymptotic state. This is the information-theoretic basis of confinement (see Section 8).

6 The n=2 Fermion: Leptons

The n=2 sector is the simplest non-trivial fermion: a two-site equal superposition with $\|F\| = \|B\| = 1/\sqrt{2}$.

6.1 Charged lepton ($m = 1$)

For $\psi = (e_0 + ie_1)/\sqrt{2}$ (winding m=1), the compression residual B is purely antisymmetric with eigenvalues $\pm 1/2$ and the universal $-\pi/2$ phase of Paper VII. The fermion has a self-contained antisymmetric residual and propagates freely. This is the charged lepton.

What Standard Model language calls a "photon" is the name given to the correlation between two such fermion events: when one charged lepton recoils, another does so in a correlated way. The codec encodes this correlation in the off-diagonal of ρ. No photon particle is required; the fermion's complexity class already contains the full correlation structure.

6.2 Neutral fermion ($m = 0$)

For $\psi = (e_0 + e_1)/\sqrt{2}$ (winding m=0), the residual B is purely symmetric with eigenvalues $\pm 1/2$ and zero net phase winding. The fermion has zero electromagnetic phase structure: it is electrically neutral. This is the neutrino.

Within the n=2 sector, the winding number acts as the charge selector:

- $m = 0$: zero phase winding $\Rightarrow$ neutral fermion (neutrino).
- $m = 1$: one phase winding $\Rightarrow$ charged fermion (electron, muon, or tau, depending on complexity level).

Electric charge is the U(1) winding number within the n=2 sector. No additional postulate is required.

7 The n=3 Fermion: Quarks and Confinement

7.1 Fractional occupancy

For a three-site equal superposition, each site carries probability $|\psi_k|^2 = 1/3$. This is the per-site occupancy of a quark in a colour triplet. The value 1/3 requires no postulate: it is the Born-rule probability of finding the fermion on any one colour site.

7.2 Colour structure ($m = 1$)

For winding m=1, the compression residual has off-diagonal entries with equal magnitude:

$$|B_{ij}| = \frac{1}{3} \quad \text{for all colour pairs } i \neq j.$$

This is exact colour symmetry, emerging from the equal-superposition constraint and Theorem 1. No SU(3) symmetry is imposed.

What Standard Model language calls a “gluon” is the name given to the correlation between quark colour changes. The mixed symmetry of B for n=3, m=1 means this correlation has no self-contained symmetry class — see Section 8.

7.3 Scalar sector ($m = 0$)

For winding m=0, the residual is purely symmetric: a spin-0, colour-neutral scalar with three-fold internal structure. Whether this corresponds to the physical Higgs boson requires derivation of the mass generation mechanism, which is open.

8 Confinement from Mixed Symmetry

The n=2, m=1 fermion (charged lepton) has a purely antisymmetric residual. Antisymmetry is self-contained: projection onto the antisymmetric subspace recovers the residual exactly, with zero error. The fermion can be described as a free asymptotic state.

The n=3, m=1 fermion (quark) has a mixed-symmetry residual. Projection onto either self-contained class (symmetric or antisymmetric) leaves a nonzero residual of magnitude ≈ 0.408, verified numerically. A mixed-symmetry fermion always requires reference to the configuration from which it arose.

Conjecture 1 (Confinement from mixed symmetry)**.** A fermion whose compression residual B has mixed symmetry cannot be described as a free asymptotic state under any minimum-description-length codec. Colour confinement is a structural consequence of the three-site codec geometry, not a separately imposed force law.

9 The Fermion Classification

Table 2 presents the fermion complexity classes. The right-hand column gives the Standard Model label for the fermion event correlations that an observer using that language would identify — not the name of a fundamental particle, but the name given to a pattern of fermion behaviour.

(n,m)	$\|F\|$	$\|B\|$	Symmetry	Fermion class	SM correlation label	Status
$(1,0)$	1	0	—	Localised, no residual	vacuum	exact
$(2,0)$	$1/\sqrt{2}$	$1/\sqrt{2}$	symmetric	Neutral lepton	neutrino	exact
$(2,1)$	$1/\sqrt{2}$	$1/\sqrt{2}$	antisymmetric	Charged lepton	electron/muon/tau	exact
$(3,0)$	$1/\sqrt{3}$	$\sqrt{2/3}$	symmetric	Colour-neutral scalar	Higgs sector	conjecture
$(3,1)$	$1/\sqrt{3}$	$\sqrt{2/3}$	mixed	Colour-charged, confined	quark/gluon	exact
$(4,2)$	$1/2$	$\sqrt{3}/2$	symmetric	Spin-2 fermion class	open	open

Table 2: Fermion complexity classes. The "SM correlation label" column gives the Standard Model name for the correlation pattern an observer would attribute to this fermion class. These are derived labels, not fundamental particles. "Status" refers to the (n,m) norm structure of Section 3, which is exact in every row; it does *not* certify that all SM particles sharing a label are distinguished by the framework. In particular the three charged leptons share the single class $(2,1)$ and are not yet distinguished by mass — see Section 10.

10 Mass and the Missing Generation Index

The spectral complexity C_s assigns an information cost proportional to the dominant wavenumber of the fermionic wavefunction. Under Solomonoff induction, $P(\psi) \propto 2^{-C_s(\psi)}$, so a genuine mass-from-complexity relation would order observable fermions by ascending C_s.

The (n,m) classification of Section 5 does not yet support this. All three charged leptons — electron, muon, tau — occupy the *same* class, $(n,m) = (2,1)$. The classification distinguishes charge and confinement (m as winding number, n as site count) but contains no third parameter that could distinguish particles of different mass within a class. A genuine derivation of the lepton mass hierarchy from C_s therefore requires extending the classification with an explicit generation index, which does not yet exist in this framework.

A combinatorial argument is sometimes proposed as a candidate origin for a discrete mass unit: one fermion hop requires four binary choices (which site, which frame, real or imaginary component, sign of phase), giving $2^4 = 16$ configurations and suggesting a 4-bit granularity. This argument is independently motivated but is *not* connected by any derivation to the C_s formula of Paper VI or to the observed mass gaps in Table 3. No construction in this framework takes $(n,m) = (2,1)$ together with a generation label and outputs $C_s = 8$ or $C_s = 12$. The near-integer appearance of

Fermion	Mass (MeV)	$\log_2(m/m_e)$	Notes
Electron	0.511	0.00	
Muon	105.66	7.69	
Tau	1776.86	11.76	

Table 3: Charged lepton masses in $\log_2(m/m_e)$ space. The two gaps are 7.69 and 4.07 bits respectively — not equal to each other, and not both close to a common integer unit. No formula in this paper computes these numbers from C_s.

$7.69 \to 8$ relies on rounding a number that is not itself derived, and the second gap (4.07) is not close to the same unit that the first gap is rounded to.

We therefore do not claim a mass ladder. The honest status is: the (n, m) classification correctly separates fermions by charge and confinement class, but says nothing yet about why three charged-lepton masses exist within the single class $(2, 1)$, or what values they should take. This is stated as an open problem in Section 12 rather than as a partial result.

11 Discussion

11.1 Why we only observe fermions

The foundational asymmetry of the codec framework is this: fermions are pixels, bosons are compression residuals. A pixel can be observed — it is a localised excitation at a definite site, sampled by the Born rule. A compression residual cannot be observed directly, because it is not a configuration of the system — it is the codec's internal record of how the configuration changed.

Consider what actually happens in a particle physics experiment. An electron enters a detector. The detector registers a hit at a definite location. Another electron, some distance away, is deflected. Standard Model language inserts a photon between these two events and says the photon "mediated the interaction." But no detector ever registers a photon as a localised pixel. What detectors register are always fermion events — photoelectric hits, track ionisations, Cherenkov rings. The photon is the name we give to the correlation between those events.

In the codec framework, that correlation is already present in the off-diagonal of the density matrix. The fermion in n=2, m=1 superposition has an antisymmetric residual B that encodes, exactly and completely, the correlation structure between its transition events. No separate photon particle is needed. The Standard Model photon is a derived label, not a fundamental object.

This is not a new claim in physics. It echoes the Wheeler–Feynman absorber theory, the S-matrix programme of the 1960s, and the modern amplitudes programme — all of which question whether the virtual particle picture is physically necessary. The codec framework makes the claim precise: virtual particles are codec artifacts, and real particles are fermions.

11.2 What the open problems reduce to

The previous version of this paper listed eight open problems, most of which concerned boson physics: W/Z mass generation, the Higgs mechanism, gauge symmetry unification, the graviton identification. Under the fermion-first reframing, these dissolve or transform.

The W and Z bosons are observed as fermion recoil events at specific energy scales. The question is not "what is the W boson?" but "why do fermion transitions at those energy scales produce correlation patterns that are massive and short-ranged?" That is a question about the lognormal probability distribution of Paper V: fermion transitions near or above the lognormal peak are expected to be suppressed and short-range. This connection is plausible but not yet quantitative, and does not rely on any claimed mass ladder.

The Higgs mechanism in the Standard Model gives mass to the W and Z by coupling them to a scalar field. In the codec framework, the prior question is more basic: why does the n=2, m=1 fermion class contain particles of more than one mass at all, when the (n, m) classification of Section 5 has no parameter to distinguish them (Section 10)? Until that generation index is identified, no statement about a Higgs-like mass generation mechanism can be derived from this framework; it remains speculative.

The graviton question is genuinely open and worth stating carefully. Gravity was derived in Papers IV–V at the geometric layer, from the aspect ratio and counting equations. The n=4, m=2 fermion class has codec spin 2. Whether these are the same physical object requires a cross-layer analysis. We note the coincidence without claiming a derivation.

11.3 The three genuine open problems

Stripped of boson physics, the genuine open problems are:

1. **What distinguishes the charged lepton generations.** The (n, m) classification places electron, muon, and tau in the same class, $(2, 1)$. No parameter in the present framework varies across them. Identifying such a parameter — and only then asking whether it explains the observed masses, and whether it caps the count at three — is the prerequisite open problem.

2. **Why quarks have the masses they have.** The six quarks span several orders of magnitude and have not been connected to C_s in any form, pending resolution of the generation-index problem above.

3. **Why the fermion force laws have the form they have.** The $1/r^2$ force law was derived in Paper VII from graviton flux conservation. The electromagnetic $1/r^2$ law should follow from the same argument applied to the n=2, m=1 fermion's antisymmetric residual. This derivation is not yet complete.

Everything else — boson masses, gauge symmetry groups, coupling constants — is a derived description of fermion correlation patterns, not a fundamental open problem.

11.4 The relationship between geometric and codec layers

Papers IV–V derived the geometric structure of spacetime from the aspect ratio and counting equations, without reference to the wavefunction codec. Papers VI–VIII derive quantum mechanics and particle structure from the wavefunction codec, without reference to the geometric layer.

Both derivations are self-consistent. The question is how they connect. The natural bridge is the worldline: a fermion's trajectory through spacetime is a complex worldline $\psi_{\text{traj}}(t) = x(t) + iy(t)$, and Paper VII showed that the minimum spectral complexity closed worldline is the ellipse — Kepler's orbit. The fermion codec and the geometric gravity are therefore not independent: the minimum-C_s worldline in the geometric spacetime is exactly the trajectory selected by the fermion's Solomonoff induction. A full derivation of this bridge is the primary target of the next paper.

12 Open Problems

1. **A generation index for charged leptons.** Identify a parameter, absent from the present (n, m) classification, that distinguishes electron, muon, and tau. Only once such a parameter exists can its connection (if any) to C_s, to the observed mass gaps, and to the number of generations be assessed.
2. **Quark mass hierarchy.** Connect the six quark masses to C_s, pending resolution of the generation-index problem above.
3. **Electromagnetic force law from fermion residual.** Apply the graviton flux argument of Paper VII to the n=2, m=1 antisymmetric residual and derive the Coulomb $1/r^2$ law.
4. **Confinement as theorem.** Prove that mixed-symmetry fermion residuals cannot be free asymptotic states under any minimum-description-length codec.
5. **Cross-layer synthesis.** Connect the geometric gravity of Papers IV–V to the fermion worldline codec. Determine whether the n=4, m=2 fermion class and the geometric graviton are the same physical entity.
6. **Many-fermion extension.** Extend the single-fermion codec to tensor products, giving composite particles and the full Fock space.

13 Summary of Results

Exact results:

1. $\|F\| = 1/\sqrt{n}$ for any n-site equal superposition.
2. $\|B\| = \sqrt{(n-1)/n}$ for any n-site equal superposition.
3. $\|F\|^2 + \|B\|^2 = 1$ for any pure state.
4. Fermion spin-1/2 is encoded in the universal $-\pi/2$ phase of Paper VII.
5. Pauli exclusion: $\|B\| = 0$ for any localised fermion.
6. Colour symmetry: $|B_{ij}| = 1/3$ for all colour pairs in the n=3, m=1 class.
7. Charge as winding number: $m = 0$ gives neutral fermion, $m = 1$ gives charged fermion, within the n=2 sector.

Conjectures (physically motivated, not yet derived):

8. Confinement from mixed symmetry of the n=3, m=1 fermion residual.
9. n=3, m=0 as the scalar sector underlying the Higgs mass generation pattern.

Open:

10. A generation index distinguishing electron, muon, and tau within the single class $(2,1)$. No such parameter exists in the present framework; the previously claimed 4-bit mass ladder is not a derived consequence of C_s and is withdrawn (see Section 10).
11. Physical identification of the n=4, m=2 spin-2 fermion class and its relation to geometric gravity.

14 Conclusion

We have proved that the density matrix of a fermion in equal superposition across n sites splits orthogonally into a fermionic diagonal F and an off-diagonal residual B satisfying

$$\|F\| = \frac{1}{\sqrt{n}}, \quad \|B\| = \sqrt{\frac{n-1}{n}}, \quad \|F\|^2 + \|B\|^2 = 1.$$

The conservation law is exact for any pure state and subsumes the Paper VII result as a special case.

The central reframing of this paper is that bosons are not fundamental particles. They are the codec's internal record of fermion transitions. Every physical observation is a fermion observation. The Standard Model boson table is a derived description of fermion correlation patterns, not an independent particle ontology.

Under this reframing, the particle classification is a *fermion* classification. The pair (n, m) describes the fermion's complexity class. The n=2 sector gives charged leptons and neutrinos, distinguished by winding number. The n=3 sector gives confined colour-charged fermions (quarks) and a scalar sector. The classification does not yet distinguish the three charged lepton generations, which all occupy the single class $(2,1)$; identifying the missing generation index is left as an open problem rather than asserted as a derived mass ladder.

The open problems reduce to three purely fermionic questions: what distinguishes the lepton generations, why the quark masses have the values they have, and whether the $1/r^2$ electromagnetic force law follows from the fermion residual structure by the same argument that gives gravity in Paper VII.

Supplementary Material

- `s1_universal_boson_theorem.py` — Theorem 1: 1000 random phase trials per n.
- `s2_spin_classification.py` — Full (n, m) classification; gluon colour structure.
- `s3_charge_and_confinement.py` — Fractional occupancy; symmetry errors; confinement projection residuals.
- `s4_mass_spectrum.py` — Reports observed lepton $\log_2(m/m_e)$ values and the combinatorial 4-bit argument side by side, without asserting a connection between them; see Section 10.
- `s5_fermion_boson_conservation.py` — Theorems 2 and 3; theta-dependent conservation; full spectrum summary.
- `mass_ladder_audit.py` — Audit the mass ladder conclusions.