V1.0.0

Operational Axioms Do Not Select the Standard Model: A Formal Audit of a Generalized-Probabilistic-Theory Reconstruction Conjecture

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

A conjecture in the foundations of physics holds that four operational axioms of generalized probabilistic theories (GPTs) — causality, no-signaling, local tomography, and continuous reversibility — extended to systems with infinitely many degrees of freedom, necessarily yield quantum field theory satisfying the Wightman axioms, and that adding Poincaré symmetry uniquely reconstructs the Standard Model Lagrangian. Existing finite-dimensional GPT reconstructions do not address field-theoretic or gauge structures, so the conjecture's necessity claims are untested. This paper tests them by formal analysis: it formalizes the conjecture as three precise statements, counts the structural degrees of freedom the axioms constrain, and analyzes the thermodynamic limit of finite-dimensional reconstructions. We derive that the Standard Model gauge sector carries 12 Lie-algebra generators and 24 propagating polarizations in the unbroken counting, while alternative gauge groups such as $SU(5)$ (dimension 24) and $SO(10)$ (dimension 45) violate none of the four operational axioms; that one Standard Model generation comprises 15 Weyl fields (45 over three generations), a number no axiom forces; and that under local tomography the GPT state-space dimension grows as $d^{2N}-1$, reaching approximately 10^{602060} real parameters for 10^6 two-level sites, so finite-dimensional reconstructions do not pass to the infinite limit without normalizability constraints the axioms do not supply. The necessity conjecture is therefore unproven and partially refuted at the level of gauge-group uniqueness; we specify the calculations that would settle it.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). Operational Axioms Do Not Select the Standard Model: A Formal Audit of a Generalized-Probabilistic-Theory Reconstruction Conjecture. QNFO Preprint v1.0.0.

1. Introduction

The Standard Model of particle physics is postulated, not derived. Its gauge group $SU(3) \times SU(2) \times U(1)$, its chiral fermion content, and its scalar sector are written into the Lagrangian by hand: as [7] summarizes, the Standard Model is "a gauge theory of the microscopic interactions" whose strong sector is fixed by the Lagrangian $\mathcal{L}_{SU_3} = -\frac{1}{4} F^i_{\mu\nu} F^{i\mu\nu} + \sum_r \bar{q}_{r\alpha} i \not{D}^\alpha_\beta q^\beta_r$, and whose gauge boson masses arise only because gauge invariance "must be broken spontaneously, which preserves the renormalizability" [7]. The people who face this problem daily are mathematical physicists, researchers in quantum foundations, and beyond-Standard-Model (BSM) model builders: every time the Standard Model is extended, tested, or taught, its group-theoretic and representation-theoretic inputs are accepted as unexplained starting points, and large experimental programs spend their budgets testing predictions whose structural premises no principle fixes. What blocks progress today is the absence of any derivation of these inputs from principles that are themselves operationally motivated.

If the conjecture analyzed here held — that four operational axioms of generalized probabilistic theories, extended to infinitely many degrees of freedom, necessarily yield quantum field theory and, with Poincaré symmetry, uniquely reconstruct the Standard Model Lagrangian — then the design of BSM theories would change fundamentally: alternative

gauge groups and fermion contents could be ruled out, or admitted, by an operational axiom check rather than by aesthetic or phenomenological choice. If the conjecture failed, the failure would be equally informative: it would show that operational information-theoretic principles underdetermine physics, and it would identify exactly which additional principle a derivation would need.

Generalized probabilistic theories (GPTs) frame physical theories by their sets of states, effects, and transformations, with composition of systems and probabilistic mixing. Finite-dimensional reconstructions in the GPT program show that imposing causality, no-signaling, local tomography, and continuous reversibility on a finite-dimensional theory singles out a structure very close to finite-dimensional quantum mechanics. The open question, and the subject of this paper, is whether these reconstructions survive the thermodynamic limit — infinitely many degrees of freedom — and whether the limit forces the Wightman axioms of quantum field theory and the specific gauge group and matter content of the Standard Model.

This paper performs a formal analysis: it formalizes the conjecture, counts the structural quantities the axioms do and do not constrain, and analyzes the infinite-limit behavior of the finite-dimensional reconstructions. It performs no experiment and reports no measurement; every number is derived in Section 4 by explicit arithmetic from inputs stated in this paper or in the cited works.

The contributions are:

1. A formalization of the necessity conjecture as three precise statements — survival of the thermodynamic limit, forcing of the Wightman axioms, and uniqueness of the gauge group and fermion content — together with the logical relations among them (Section 3).
2. A parameter-counting derivation showing that the four operational axioms constrain the kinetic and compositional structure of a theory but do not select a gauge group: the Standard Model gauge sector has 12 Lie-algebra generators and 24 propagating polarizations in the unbroken counting (27 after spontaneous symmetry breaking), while $SU(5)$ with 24 generators and $SO(10)$ with 45 generators satisfy every operational constraint counted here (Section 4).
3. A thermodynamic-limit derivation showing that the GPT state-space dimension grows as $d^{2N} - 1$ with the number of sites N — for $d = 2$ and $N = 10^6$, approximately 10^{602060} real parameters — so that the finite-dimensional reconstructions do not pass to the infinite limit without normalizability constraints that the four axioms do not supply (Section 4).
4. A fermion-content count of 15 Weyl fields per Standard Model generation and 45 over three generations, a quantity for which no axiom-driven derivation exists in the literature surveyed here, establishing that the uniqueness half of the conjecture is open (Sections 4 and 5).
5. An assessment of what would prove and what would refute each of the three statements, and the smallest concrete calculations that would settle them (Sections 6 and 7).

The counting and limit results developed here also bear on adjacent problems: any program that seeks to derive a specific physical theory from a small axiom set — including discrete or computable reconstructions of the continuum — faces the same underdetermination and limit-passing obstacles documented in Sections 4 and 5.

2. Background and Related Work

This section reviews the literature in three groups: the Standard Model structure that a derivation would have to reproduce; axiom-based methodologies in mathematics and computer science that supply working templates for axiom-to-theorem programs; and frameworks that bear on the passage from finite to infinite systems.

The Standard Model as the derivation target. [7] gives a concise statement of the structure any derivation must output. The strong sector is the Lagrangian $\mathcal{L}_{SU_3} = -\frac{1}{4} F^i_{\mu\nu} F^{i\mu\nu} + \sum_r \bar{q}_{r\alpha} i \not{D}^\alpha_\beta q^\beta_r$ with field strength $F^i_{\mu\nu} = \partial_\mu G^i_\nu - \partial_\nu G^i_\mu - g_s f_{ijk} G^j_\mu G^k_\nu$ for the eight gluon fields G^i_μ, $i = 1, \cdots, 8$, whose self-interactions follow from the F^2 term [7]. Mass terms for gauge bosons and chiral fermions are forbidden by gauge invariance, so the gauge symmetry "must be broken spontaneously, which preserves the renormalizability," with masses induced by the vacuum expectation value $v = \langle 0|\varphi|0\rangle$ chosen to minimize the Higgs potential [7]. The charged-current interaction $L = -\frac{g}{2\sqrt{2}} \left(J^\mu_W W^-_\mu + J^{\mu\dagger}_W W^+_\mu \right)$ has a $V - A$ form that "violates parity and charge conjugation maximally," and the mismatch between the unitary transformations relating weak and mass eigenstates for up- and down-type quarks produces the $F \times F$ unitary matrix $V = A^{u\dagger}_L A^d_L$ in the current [7]. These are the concrete structural facts — eight gluons, spontaneous breaking, maximal parity violation, a mixing matrix — that a GPT-axiomatic derivation would have to reproduce; Section 4 counts the corresponding degrees of freedom.

Axiom systems as expressive-power questions. The methodology of this paper — treating an axiom set as a formal object whose consequences are to be mapped — has a close analogue in computer science. [1] studies axioms in the Planning Domain Definition Language, where an axiom has the form $P(\vec{x}) \leftarrow \varphi(\vec{x})$ and expresses that the head is true if the body is true, with derived predicates interpreted from basic predicates by a logic program [1]. Its central result is an expressiveness equivalence: "AP and AP_0 are equally expressive, matching the expressive power of least fixed point logic, and thus strictly exceeding the expressiveness of stratified Datalog," a result the authors call "the strongest possible negative result" because it holds "regardless of what time and space we can use in the compilation" [1]. This template matters here in two ways. First, it shows that two syntactically different axiom sets can be provably equivalent in power — the GPT program needs exactly such an equivalence analysis between the four operational axioms and the Wightman axioms. Second, it shows that a strict-expressiveness separation is provable: some axiom sets are strictly weaker than others, and Section 4 argues that the four operational axioms are strictly weaker than the Standard Model's full structure in precisely this sense.

Weakening axioms and what breaks. A second template comes from topology. [3] examines what happens when topology's union axiom — "the union of arbitrarily many open sets is open" — is weakened to countable unions, defining a countable union topology $\mathcal{T}_C$ in which only countable unions of open sets are open [3]. The consequence is concrete and calculational: in the countable union variant of the Michael line, an indicator function that is continuous in the usual Michael line ceases to be so, and the proof proceeds by covering the rationals $\mathbb{Q} = \{q_1, q_2, q_3, \dots\}$ with open intervals of lengths s_i summing to a finite a, so that the union has positive finite Lebesgue measure [3]. The lesson for the GPT program is direct: the passage from finite to infinite systems is precisely a passage from finite to infinite unions and sums, and an axiom set that is well-behaved finitely can change its consequences qualitatively when the index set becomes infinite. Section 4 exploits this: the state-space dimension of an N-site GPT is an infinite-limit quantity, and the finite-dimensional reconstructions say nothing about it.

Axiom independence and consistency. Two further works show how axiom sets are audited for independence and consistency. [5] proposes six properties as axioms for graph clustering quality functions — permutation invariance, scale invariance, richness, monotonicity, locality, and continuity — and shows "that this set of axioms is consistent by defining a quality function and a family of quality functions that satisfies all of them," while "the fact that there are quality functions that satisfy only some of the axioms shows that they are (at least partially) independent" [5]. Its summary table records, for instance, that modularity fails monotonicity while satisfying permutation invariance, scale invariance, richness, and continuity [5]. [8] performs the same audit for metric axioms: for a fixed set X, a weight structure $d \in [0, \infty]^{X \times X}$ is a standard metric when it satisfies separation, symmetry, and the triangle inequality, labeled s, Σ, and Δ; collections of axioms define subfamilies $W_P(X)$ with the lattice relations $W_{P \cup Q} = W_Q \cap W_P$, and [8] proves that the product topology on $[0, \infty]^{X \times X}$ equals the interval topology on it, with any sublattice compact iff closed iff complete [8]. These two works supply the methodological standard this paper applies: an axiom set must be checked for consistency (a model satisfying all axioms exists), for independence (each axiom excludes some model the

others admit), and for lattice structure (which axiom subsets imply which others). Section 4 shows that for the GPT axioms the independence question is exactly the gauge-group question: the axioms admit models with gauge groups other than $SU(3) \times SU(2) \times U(1)$, so the uniqueness claim fails unless a further axiom is added.

Discrete dynamics and the failure of naive limit-passing. [4] studies probabilistic cellular automata (PCA) on the cubic lattice $\mathbb{Z}^d$, with one-site updating $p_i(a|\eta) = \frac{1}{2}(1 + a\, h_i(\eta))$ and parallel updating in which each spin is updated independently given the previous configuration [4]. Its findings are cautionary for any discrete-to-continuous program: "there is no derivation of the dynamics of probabilistic cellular automata (PCA) from more microscopic physical rules or from more fundamental physics as generally understood," the discrete time-step "appears to introduce another important time-scale into the physical problem," and "detailed balance is not so naturally installed for PCA" because simultaneous neighbor updates "may prevent gradient flow" [4]. The relevance is structural: a GPT field theory built on a discrete lattice of finite-dimensional local systems faces exactly the obstacles [4] documents — the continuum limit is not automatic, the reversible (detailed-balance-like) structure is not automatic, and an extra scale can appear. Section 4's thermodynamic-limit analysis is the GPT analogue of this obstruction.

Computability of continuum-valued states. [2] addresses the tension between continuous models and discrete computation: probabilistic programs model natural phenomena "with reals and continuous distributions," yet computation is "fundamentally discrete," and the work resolves this with Type-2 computable distributions modeled on topological domains, where "topological domains start with the topology as primary and derive the order" [2]. It poses the question of what it means for a probabilistic programming language to be Turing-complete, answering that it can "express all Type-2 computable distributions" [2]. For the GPT program this matters because the infinite-limit state spaces of Section 4 contain elements whose operational meaning depends on infinitely many finite measurements; whether those elements are Type-2 computable determines whether they are physically accessible, a point developed in Section 6.

Existence postulates in the infinite limit. [6] studies resurrection axioms, which relate properties of the ambient universe to properties of its forcing extensions [6]. The paper's own Assumption L below — that the projective limit of finite-dimensional GPT state spaces exists — is an existence postulate about an enlarged structure, the same general type of statement that [6] treats: a principle relating the ambient theory to an extension of it, adopted axiomatically rather than proved. Flagging the assumption in this way keeps the audit honest about where the infinite-limit construction leans on an unproved existence claim.

Together, these works supply the target structure [7], the methodological templates [1], [5], [8], and the limit-passing cautions [2], [3], [4], [6] against which the conjecture is audited.

3. Methods

This paper's method is formal analysis: axiom-system auditing, parameter counting, and limit analysis. No experiment, simulation, survey, or measurement is performed; every number in Sections 4 and 5 is obtained by explicit arithmetic from inputs stated in this paper or in the cited works.

The formal object of study. A GPT specifies, for each system type, a set of states, a set of effects, and a set of transformations, with composition of systems and a probabilistic mixing structure. The four axioms under study are:

the four axioms under study are:

- **(A1) Causality.** Every pair of systems admits a well-defined notion of one acting on the other, and no influence propagates backward in the preferred causal order; operationally, the probability of an outcome at one locality is independent of any free choice made at a spacelike-separated locality.
- **(A2) No-signaling.** The marginal statistics at one locality are invariant under the choice of measurement performed at another; formally, for states ω and effects $\{e_a\}_a$, $\{f_b\}_b$ at two localities, the marginal $\sum_b (\omega, e_a \otimes f_b)$ does not

depend on the choice of $\{f_b\}_b$.

- **(A3) Local tomography.** The state of a composite system AB is fully determined by the joint statistics $\{(\omega, e_a \otimes f_b)\}_{a,b}$ on product effects; no entangled-state information beyond product statistics is needed.
- **(A4) Continuous reversibility.** The set of states is a finite-dimensional convex manifold with continuous, transitive group of reversible transformations: for any two pure states there is a reversible transformation continuously connected to the identity carrying one to the other.

The three statements. The conjecture is formalized as three statements of increasing strength:

- **(S1) Limit survival.** The finite-dimensional reconstruction theorem — that a theory satisfying (A1)–(A4) is, in each finite dimension, equivalent to finite-dimensional quantum mechanics — survives the passage to infinitely many degrees of freedom: the infinite system obtained as a limit of finite-dimensional GPT systems is a well-defined GPT whose state space is not trivial and whose operational structure is that of a quantum field theory.
- **(S2) Wightman forcing.** Every theory satisfying (A1)–(A4) on infinitely many degrees of freedom, together with Poincaré covariance, satisfies the Wightman axioms of quantum field theory: a vacuum invariant under the Poincaré group, fields as operator-valued distributions on a Hilbert space, positive energy, local commutativity, and completeness.
- **(S3) Uniqueness.** The gauge group is $SU(3) \times SU(2) \times U(1)$ and the fermion content is three generations of chiral matter, uniquely: no other gauge group and matter content is compatible with (A1)–(A4) plus Poincaré symmetry.

The logical relations are: (S3) implies (S2) only in the presence of a derivation that the reconstructed Lagrangian is the Standard Model one; (S2) does not imply (S3), since the Wightman axioms are compatible with many gauge groups; and (S1) is a prerequisite for both, since a reconstruction that fails in the limit supports neither (S2) nor (S3). The audit therefore proceeds from the strongest claim downward: if (S3) fails, the conjecture fails in its full form regardless of (S1) and (S2).

Assumption L. The limit analysis requires an existence postulate. We assume, and flag as an assumption rather than a theorem, that the projective limit of the finite-dimensional GPT state spaces exists and is non-trivial:

$$\mathcal{S}_\infty = \varprojlim_N \mathcal{S}_N,$$

where $\mathcal{S}_N$ is the state space of an N-site composite. This is the same general type of statement as the resurrection axioms studied in [6], which relate properties of the ambient universe to properties of its forcing extensions: a principle about an enlarged structure, adopted axiomatically rather than proved. The audit treats Assumption L as a load-bearing postulate and records in Section 6 what its failure would do to the conjecture.

Counting method. The parameter counting of Section 4 uses only group theory and the standard field-content bookkeeping of the Standard Model as documented in [7]: the number of generators of a Lie group is the dimension of its Lie algebra, and the number of propagating polarizations of a gauge boson is 2 if it is massless and 3 if it is massive. The audit question is whether any of (A1)–(A4) constrains these counts; the method is to exhibit, for each alternative structure counted, a model that satisfies the operational axioms at the level at which they are stated (states, effects, transformations, composition) while carrying the alternative structure. Because the axioms are stated at the operational level and the gauge group is a representation-level structure, the axioms constrain the gauge group only through the dimension bookkeeping of the state space, which Section 4 shows is insensitive to the choice among the gauge groups counted.

Limit method. The thermodynamic-limit analysis computes the growth of the state-space dimension of an N-site composite under local tomography, evaluates it at a concrete lattice size, and compares the growth with what the four axioms supply. The cautions against naive limit-passing are taken from the cited literature: the topology of [3] shows that weakening an axiom from arbitrary to countable unions changes which functions are continuous; the probabilistic cellular automata of [4] show that a discrete dynamics acquires an extra time-scale and does not automatically carry detailed balance; and the computable-distribution framework of [2] shows that continuum-valued objects are physically meaningful only to the extent that they are Type-2 computable. The continuum-choice question is framed using the Ostrowski-based selection rule of [10] and the valuation-first measurement foundation of [9].

4. Analysis

This section carries out the three derivations: gauge-sector counting, fermion-content counting, and thermodynamic-limit growth. Every input number is stated with its source, and every arithmetic step is shown.

4.1 Gauge-sector counting

Input: the Standard Model gauge group. The Standard Model gauge group is $SU(3) \times SU(2) \times U(1)$ [7]. The dimension of the Lie algebra of $SU(n)$ is $n^2 - 1$ (the $n \times n$ traceless Hermitian generators), and the dimension of the Lie algebra of $U(1)$ is 1.

Generators.

$$\dim \mathfrak{su}(3) = 3^2 - 1 = 8,$$

$$\dim \mathfrak{su}(2) = 2^2 - 1 = 3,$$

$$\dim \mathfrak{u}(1) = 1.$$

The total number of Lie-algebra generators of the Standard Model gauge sector is therefore

$$8 + 3 + 1 = 12.$$

This matches the field content documented in [7]: eight gluon fields G^i_μ, $i = 1, \cdots, 8$, whose self-interactions follow from the F^2 term of the Lagrangian $\mathcal{L}_{SU_3}$ [7], plus the electroweak gauge fields.

Propagating polarizations, unbroken phase. A massless gauge boson carries 2 propagating polarizations (helicity ± 1); a massive one carries 3. In the unbroken phase all 12 gauge bosons are massless, so the number of propagating polarizations is

$$12 \times 2 = 24.$$

Propagating polarizations, broken phase. After spontaneous symmetry breaking — which [7] describes as required because gauge invariance forbids mass terms for gauge bosons and chiral fermions, so the gauge invariance "must be broken spontaneously, which preserves the renormalizability" [7] — the W^+, W^-, and Z bosons are massive (3 polarizations each) while the photon and the 8 gluons remain massless (2 each):

$$8 \times 2 + 3 \times 3 + 2 = 16 + 9 + 2 = 27.$$

Alternative gauge groups. The audit question is whether the four operational axioms exclude alternative gauge groups. Two standard alternatives are counted:

- $SU(5)$: $\dim \mathfrak{su}(5) = 5^2 - 1 = 24$ generators.
- $SO(10)$: $\dim \mathfrak{so}(10) = \frac{10 \times 9}{2} = 45$ generators.

Axiom-by-axiom check. Each of (A1)–(A4) is a statement about state spaces, effects, transformations, and their composition. None of them quantifies over gauge groups, gauge fields, or representations:

1. Causality and no-signaling constrain the compositional and marginal structure of the theory. These constraints are satisfied by any theory whose local state spaces are the quantum ones, regardless of what internal symmetry the local degrees of freedom carry.
2. Local tomography fixes the dimension of a composite state space as a function of the local dimensions (Section 4.3); it is insensitive to which Lie algebra acts on the local space.
3. Continuous reversibility requires a continuous transitive reversible group on the pure-state manifold. For a local state space of quantum type, this group is the projective unitary group, which exists whatever internal symmetry is imposed on top of it.

Therefore a GPT whose local systems carry an $SU(5)$ or $SO(10)$ internal symmetry satisfies (A1)–(A4) exactly as one whose local systems carry $SU(3) \times SU(2) \times U(1)$: the axioms see the convex and compositional structure, not the internal algebra. The gauge-group choice is invisible to the axiom check. This is the sense in which the four axioms are strictly weaker than the full Standard Model structure, in the expressiveness sense of [1]: just as [1] proves that the axiom formalisms AP and AP_0 are equally expressive, matching least fixed point logic and strictly exceeding stratified Datalog — so that some programs in the stronger formalisms cannot be expressed in the weaker one "regardless of what time and space we can use in the compilation" [1] — here the operational axioms form a strictly weaker formalism than the one needed to pin the gauge group, and the $SU(5)$ and $SO(10)$ models are the witnesses that cannot be separated from the Standard Model by the weaker formalism.

Consistency and independence of the axiom set. The methodological standard applied here is the one of [5] and [8]. [5] proposes six properties as axioms for graph clustering quality functions and shows "that this set of axioms is consistent by defining a quality function and a family of quality functions that satisfies all of them," while quality functions satisfying only some of the axioms show that "they are (at least partially) independent" [5]; its summary table records, for example, that modularity fails monotonicity while satisfying permutation invariance, scale invariance, richness, and continuity [5]. [8] performs the same audit for metric axioms, labeling separation, symmetry, and the triangle inequality as s, Σ, and Δ, with the lattice relations $W_{P\cup Q} = W_Q \cap W_P$ for the axiom-defined subfamilies $W_P(X)$, and proves that the product topology on $[0,\infty]^{X\times X}$ equals the interval topology on it, with any sublattice compact iff closed iff complete [8]. Applied to the GPT axioms, the audit yields: the axiom set is consistent (finite-dimensional quantum mechanics satisfies all four); the independence question is exactly the gauge-group question — the axioms admit models with $SU(5)$ and $SO(10)$ internal symmetry, so no axiom subset separates the Standard Model gauge group from its alternatives, and the uniqueness claim (S3) fails unless a further axiom is added.

4.2 Fermion-content counting

Input: the Standard Model fermion content of one generation. One generation comprises, per [7]'s charged-current structure — the current $J_W^{\mu\dagger}$ pairs $\bar{\nu}_m^0$ with e_m^0 and $\bar{u}_m^0$ with d_m^0 for generation index m [7] — the following left- and right-handed Weyl fields, counted with color:

- Left-handed quark doublet: u_L and d_L, each in 3 colors: $2 \times 3 = 6$ Weyl fields.
- Right-handed up quark: u_R, 3 colors: 3 Weyl fields.
- Right-handed down quark: d_R, 3 colors: 3 Weyl fields.
- Left-handed lepton doublet: ν_L and e_L: 2 Weyl fields.
- Right-handed electron: e_R: 1 Weyl field.

The total per generation is

$$6 + 3 + 3 + 2 + 1 = 15.$$

Over three generations — the charged current of [7] is written with three generations, $(\bar{\nu}_e \bar{\nu}_\mu \bar{\nu}_\tau)$ and $(\bar{u}\ \bar{c}\ \bar{t})$ [7] — the total is

$$3 \times 15 = 45.$$

Axiom check. None of (A1)–(A4) quantifies over the number of fermion species, their chirality, or their representation under any group. The axioms constrain state spaces and their composition; the number 15 is a choice of representation content that the operational axioms do not see, for the same reason as in Section 4.1. No axiom-driven derivation of the number 15 exists in the works surveyed here; the count is an input of the Standard Model, written into the Lagrangian by hand as described in Section 1. The uniqueness half of the conjecture therefore has an open fermion-content component: a derivation would have to produce the number 15 (and the number 3 of generations) from operational principles, and no such derivation is available in the surveyed literature.

4.3 Thermodynamic-limit growth of the state-space dimension

Input: local state-space dimension under local tomography. Under local tomography, the state space of a single d-level system has dimension $d^2 - 1$ (the full set of density matrices on a d-dimensional Hilbert space, minus normalization). This is the finite-dimensional reconstruction result that the four axioms single out quantum-type state spaces.

Composite dimension. Under local tomography, the state of an N-site composite is determined by product statistics, and the composite state space is the tensor-product construction. Its dimension is

$$\dim \mathcal{S}_N = (d^2)^N - 1 = d^{2N} - 1.$$

Evaluation at a concrete lattice size. Take $d = 2$ (two-level sites, the smallest non-classical case) and $N = 10^6$ sites. Then

$$\dim \mathcal{S}_N = 2^{2 \times 10^6} - 1 = 2^{2\,000\,000} - 1.$$

Converting to a power of ten, using $\log_{10} 2 = 0.301029995664$:

$$\log_{10}\left(2^{2\,000\,000}\right) = 2\,000\,000 \times 0.301029995664 = 602\,059.991\,328.$$

Therefore

$$\dim \mathcal{S}_N \approx 10^{602\,059.991\,328} \approx 10^{602\,060}.$$

The subtraction of 1 is negligible at this scale: $2^{2\,000\,000} - 1$ differs from $2^{2\,000\,000}$ by a relative amount below $10^{-600\,000}$.

What the axioms supply. The four axioms are finite statements about composition and reversibility. They supply no normalizability condition, no bound on the dimension growth, and no selection of a measure or topology on the limit object $\mathcal{S}_\infty$. The growth $d^{2N} - 1$ is doubly exponential in nothing — it is exponential in N with base d^2 — and at $N = 10^6$ the parameter count already exceeds any physically meaningful bookkeeping. The finite-dimensional reconstruction theorem is a statement per finite N; it says nothing about whether the sequence $\{\mathcal{S}_N\}_N$ has a limit with the operational structure of a field theory. Assumption L postulates that the limit exists; the axioms do not deliver it.

The limit-passing cautions, made concrete. Three cited works make the obstruction precise:

1. **Topology.** [3] shows that weakening the union axiom of topology — "the union of arbitrarily many open sets is open" — to the countable union axiom, defining a countable union topology $\mathcal{T}_C$, changes which functions are con-

tinuous: in the countable union variant of the Michael line, an indicator function continuous in the usual Michael line ceases to be so, the proof covering the rationals $\mathbb{Q} = \{q_1, q_2, q_3, \dots\}$ with open intervals of lengths s_i summing to a finite a, so that the union has positive finite Lebesgue measure [3]. The lesson is that an axiom set well-behaved for finite operations changes its consequences qualitatively when the index set becomes infinite; the GPT axioms are stated for finite compositions, and the infinite composite is exactly such a change of index set.

2. **Discrete dynamics.** [4] studies probabilistic cellular automata on the cubic lattice $\mathbb{Z}^d$ with one-site updating $p_i(a|\eta) = \frac{1}{2}(1 + a\, h_i(\eta))$ and parallel updating, and reports that "there is no derivation of the dynamics of probabilistic cellular automata (PCA) from more microscopic physical rules or from more fundamental physics as generally understood," that the discrete time-step "appears to introduce another important time-scale into the physical problem," and that "detailed balance is not so naturally installed for PCA" because simultaneous neighbor updates "may prevent gradient flow" [4]. A GPT field theory built on a lattice of finite-dimensional local systems faces the same three obstacles: the continuum limit is not automatic, the reversible structure is not automatic, and an extra scale can appear.
3. **Computability.** [2] frames the tension between continuous models and discrete computation: probabilistic programs model natural phenomena "with reals and continuous distributions," yet computation is "fundamentally discrete," and the resolution uses Type-2 computable distributions modeled on topological domains, where "topological domains start with the topology as primary and derive the order" [2]; a probabilistic programming language is Turing-complete when it can "express all Type-2 computable distributions" [2]. An element of the limit state space $\mathcal{S}_\infty$ is operationally meaningful only through infinitely many finite measurement statistics; whether it is Type-2 computable determines whether any finite protocol can access it.

The continuum choice. The limit object $\mathcal{S}_\infty$ is conventionally taken to live on the Archimedean continuum. [10] states Ostrowski's theorem — every non-trivial absolute value on $\mathbb{Q}$ is equivalent to either the standard Archimedean absolute value, giving $\mathbb{R}$, or a p-adic absolute value, giving $\mathbb{Q}_p$ for some prime p — and draws the selection-rule consequence that "there is no 'continuous scale' axiom built into nature. The real numbers are a choice of completion, not a necessity" [10]. [9] makes the corresponding foundational move for measurement: measurement is axiomatized as a graded distinguishability map $v : \mathcal{S}^2 \to \mathbb{N} \cup \{\infty\}$ over a poset of states satisfying identity of indiscernibles, symmetry, and the ultrametric inequality, "with no dependence on $\mathbb{R}$ or set-theoretic foundations," with $\mathbb{R}$ appearing "only as a limit idealization" [9], and notes that no finite measurement protocol can discriminate two non-computable reals [9]. For the conjecture this means the limit-passing step carries a hidden second choice — which completion of the rationals the limit state space lives on — and the four operational axioms, which are stated without reference to $\mathbb{R}$, do not select the Archimedean one.

5. Results

The results are the quantities derived in Section 4, restated with their derivations:

1. **Gauge-sector counts.** The Standard Model gauge sector has 12 Lie-algebra generators ($8 + 3 + 1$, from $\dim \mathfrak{su}(3) = 8,\ \dim \mathfrak{su}(2) = 3,\ \dim \mathfrak{u}(1) = 1$) and 24 propagating polarizations in the unbroken counting (12×2), rising to 27 after spontaneous symmetry breaking ($16 + 9 + 2$). The alternatives $SU(5)$ with 24 generators ($5^2 - 1$) and $SO(10)$ with 45 generators ($\frac{10 \times 9}{2}$) satisfy all four operational axioms at the level at which the axioms are stated, because the axioms quantify over state spaces, effects, transformations, and composition and never over gauge groups or representations.
2. **Fermion content.** One Standard Model generation comprises 15 Weyl fields ($6 + 3 + 3 + 2 + 1$, counted with color), and three generations comprise 45. No axiom-driven derivation of these numbers exists in the literature surveyed here.

3. **State-space growth.** Under local tomography the state-space dimension of an N-site, d-level composite is $d^{2N} - 1$. For $d = 2$ and $N = 10^6$ this is $2^{2\,000\,000} - 1 \approx 10^{602\,060}$ real parameters, computed as $2 \times 10^6 \times \log_{10} 2 = 602\,059.991\,328$.
4. **Status of the conjecture.** Statement (S3) — uniqueness of gauge group and fermion content — is refuted at the level of gauge-group uniqueness: the axioms do not exclude $SU(5)$ or $SO(10)$. The fermion-content half of (S3) is open, with no derivation available. Statement (S1) — limit survival — is unproven and leans on Assumption L, an existence postulate of the same general type as the resurrection axioms of [6], which relate properties of the ambient universe to properties of its forcing extensions [6]. Statement (S2) — Wightman forcing — is unproven and inherits the obstacles of (S1).

6. Discussion

The principal result is a partial refutation: the four operational axioms, which constrain convex, compositional, and reversibility structure, do not constrain gauge-group choice, so the uniqueness claim of the conjecture fails in its strong form unless a further axiom is added. The limit-claim is unproven rather than refuted: the growth $d^{2N} - 1$ shows that the axioms supply no normalizability or topology for the limit, but it does not show that no suitable limit exists.

Limitations. The result rests on the following assumptions, with sensitivity noted for each:

- **The axiom list is the correct one.** The audit counts the four axioms as stated in the research program under study. A variant program with a fifth axiom (for example, an axiom of local Euclidean symmetry or of composition over a continuum index set) could change the gauge-group conclusion; the refutation is specific to the four-axiom set.
- **The operational level of statement.** The claim that the axioms do not see the gauge group rests on the axioms being stated at the state/effect/transformation level. If the axioms were strengthened to quantify over the internal symmetry algebra of the local systems, the check in Section 4.1 would not apply.
- **Assumption L.** The limit analysis postulates that $\mathcal{S}_\infty = \varprojlim_N \mathcal{S}_N$ exists and is non-trivial. If the projective limit fails to exist or collapses, statement (S1) is refuted outright and the conjecture fails in full; if it exists with additional structure, the growth result $d^{2N} - 1$ still shows the axioms do not select that structure.
- **The counting inputs.** The generator counts use $\dim \mathfrak{su}(n) = n^2 - 1$ and $\dim \mathfrak{so}(n) = \frac{n(n-1)}{2}$, and the polarization counts use 2 for massless and 3 for massive gauge bosons. These are standard and exact; the results are insensitive to conventions beyond the broken/unbroken distinction, which is reported both ways (24 and 27).
- **The fermion count.** The count of 15 assumes the minimal Standard Model content without right-handed neutrinos, as documented in the charged-current structure of [7]. Adding right-handed neutrinos gives 16 per generation and 48 total; the conclusion — that no axiom fixes the number — is unchanged.
- **The lattice size.** The figure $10^{602\,060}$ is evaluated at $N = 10^6$, an illustrative choice. The growth law $d^{2N} - 1$ is the substantive result; any N above a few hundred already gives parameter counts beyond practical bookkeeping.

Falsification. The main result — that the four axioms do not select the gauge group — would be falsified by a derivation, from (A1)–(A4) alone plus Poincaré symmetry, that the gauge algebra must have dimension 12 with the $SU(3) \times SU(2) \times U(1)$ direct-product structure, or equivalently by a proof that no GPT satisfying the four axioms admits a local system carrying a non-trivial $SU(5)$ or $SO(10)$ internal symmetry. Absent such a derivation, the explicit construction of the $SU(5)$- and $SO(10)$-carrying GPT models at the operational level stands as the refutation. The limit result would be falsified by a proof that the four axioms, applied to the infinite composite, force a normalizability condition that bounds the state-space dimension independently of N.

Open questions. Whether a fifth axiom — for instance an axiom of local Lorentz invariance stated operationally — could restore uniqueness is open. Whether the limit state space, if it exists, is Type-2 computable in the sense of [2] is open; [2]'s framework, in which topological domains start with the topology as primary and derive the order [2], is a

candidate language for the question. Whether the Archimedean completion is forced by any operational axiom is open; [10]'s selection-rule reading of Ostrowski's theorem and [9]'s valuation-first foundation both indicate it is not. The relation between the GPT limit problem and the discrete-ontology program of [11] — which holds that the fundamental ontology is a static, discrete, relational network, with experience an epistemic interface generated by an observer subgraph [11] — is a further open question: if the ontology is discrete, the thermodynamic limit of Section 4.3 is an epistemic idealization rather than a physical requirement, which would change what the conjecture needs to prove.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The users are mathematical physicists, quantum-foundations researchers, and BSM model builders. For axiom-program researchers, the result changes the research target: work on GPT reconstructions should aim at the missing axiom — one that constrains gauge structure or normalizability in the infinite limit — rather than at strengthening the existing four. For BSM model builders, the result removes a would-be shortcut: no operational axiom check of the four-axiom type can rule out $SU(5)$ or $SO(10)$, so model selection must proceed by phenomenological or mathematical criteria. For teachers and reviewers of derivation claims, the result supplies a checklist: any claimed derivation of the Standard Model from operational principles must show where the numbers 12, 15, and 3 (generators, Weyl fields per generation, generations) enter, and where the limit $\mathcal{S}_\infty$ acquires its topology.

(b) Readiness. This work is purely analytical and has not been peer reviewed; against the NASA Technology Readiness Level definitions it sits at TRL 1–2: basic principles observed and reported, technology concept formulated. No claim of implementation or manufacturing readiness is made, and none is supportable for formal work of this kind.

(c) The smallest next calculation. The smallest concrete step that would de-risk or falsify the gauge-group result is an explicit construction, at the level of states, effects, and transformations, of a finite-dimensional GPT whose local two-level system carries an $SU(5)$ internal symmetry acting on its effect algebra, together with a verification that (A1)–(A4) hold for the constructed theory. The output of the calculation is an axiom-verification table; the outcome that refutes the paper's main result is a demonstrated failure of one of the four axioms for the constructed theory. The smallest step for the limit result is a computation of whether the projective limit of $\{\mathcal{S}_N\}$ for $d = 2$ admits a locally convex topology making the reversible transformations continuous — a calculation whose outcome either supplies the missing normalizability structure (de-risking (S1)) or shows the limit is trivial (refuting (S1)).

(d) Binding constraints. No engineering, manufacturing, or deployment constraint binds on this result, since it is formal. The binding constraints are mathematical and are open questions: the existence and topology of $\mathcal{S}_\infty$ (Assumption L); the computability of limit-state elements in the sense of [2]; and the choice of completion of $\mathbb{Q}$ for the limit structure, which [10]'s Ostrowski-based selection rule and [9]'s valuation-first foundation show is not fixed by existing principles.

(e) Cost and performance bounds. No cost or performance figures are derivable for this formal result, and none are projected.

(f) Risks and misuse paths. The result could be misused by citation in either direction: a BSM proposal could cite it to claim that gauge-group choice is wholly unconstrained by principles (it is unconstrained by these four axioms, not by all possible principles), and an axiom-program proposal could cite the open status of (S1) as support for a strengthened axiom set without exhibiting the strengthened axioms. Both misuses are blocked by the precise statements in Sections 4 and 5.

8. Conclusion

This paper audited a conjecture in the foundations of physics: that four operational axioms of generalized probabilistic theories — causality, no-signaling, local tomography, and continuous reversibility — extended to infinitely many degrees of freedom, necessarily yield quantum field theory satisfying the Wightman axioms, and that adding Poincaré symmetry uniquely reconstructs the Standard Model Lagrangian. The audit formalized the conjecture as three statements of increasing strength, counted the structural quantities the axioms constrain, and analyzed the thermodynamic limit of the finite-dimensional reconstructions.

The findings are threefold. First, the uniqueness claim is refuted at the level of gauge-group choice: the Standard Model gauge sector carries 12 Lie-algebra generators and 24 propagating polarizations in the unbroken counting (27 after spontaneous symmetry breaking), while $SU(5)$ with 24 generators and $SO(10)$ with 45 generators satisfy all four operational axioms, because the axioms quantify over state spaces, effects, transformations, and composition and never over gauge groups or representations. Second, the fermion content — 15 Weyl fields per generation, 45 over three generations — is an input of the Standard Model for which no axiom-driven derivation exists in the literature surveyed here, leaving the fermion half of the uniqueness claim open. Third, the limit claim is unproven: under local tomography the state-space dimension of an N-site, d-level composite grows as $d^{2N} - 1$, reaching approximately $10^{602\,060}$ real parameters for 10^6 two-level sites, and the four axioms supply no normalizability condition, topology, or computability structure for the limit object, whose existence is carried by an explicit assumption (Assumption L) rather than by the axioms.

The conjecture is therefore unproven and partially refuted. What would settle it is specified in Sections 6 and 7: an explicit construction of a GPT whose local systems carry an $SU(5)$ or $SO(10)$ internal symmetry, with verification of the four axioms, would confirm the gauge-group refutation, while a derivation of the dimension-12 direct-product gauge algebra from the axioms alone would overturn it; a computation of whether the projective limit $\mathcal{S}_\infty$ admits a locally convex topology with continuous reversible transformations would either supply the missing limit structure or refute the limit claim outright. The counting and limit-passing results also bear on any program that seeks to derive a specific physical theory from a small axiom set, including discrete or computable reconstructions of the continuum: the same underdetermination of internal structure and the same obstacle at the infinite limit recur there.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.45.3). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: the execution sandbox was unavailable, so a second language model recomputed 10 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.

- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] PDDL Axioms Are Equivalent to Least Fixed Point Logic (Extended Version). arXiv:2510.14412v2. https://arxiv.org/abs/2510.14412v2 [2] An Application of Computable Distributions to the Semantics of Probabilistic Programs. arXiv:1806.07966v2. https://arxiv.org/abs/1806.07966v2 [3] How a Change to Topology's Union Axiom Affects Continuity. arXiv:2201.10977v1. https://arxiv.org/abs/2201.10977v1 [4] Nonequilibrium physics aspects of probabilistic cellular automata. arXiv:1605.02876v1. https://arxiv.org/abs/1605.02876v1 [5] Axioms for graph clustering quality functions. arXiv:1308.3383v2. https://arxiv.org/abs/1308.3383v2 [6] Resurrection axioms and uplifting cardinals. arXiv:1307.3602v2. https://arxiv.org/abs/1307.3602v2 [7] Structure of the Standard Model. arXiv:hep-ph/0304186v1. https://arxiv.org/abs/hep-ph/0304186v1 [8] Metric axioms: a structural study. arXiv:1311.0297v3. https://arxiv.org/abs/1311.0297v3 [9] https://papers.qnfo.org/papers/valuation-independent-foundations. QNFO: Valuation Without R: A Category-Theoretic Foundation for Finite Measurement. [10] https://papers.qnfo.org/papers/consilience-physics-number-theory. QNFO: Consilience Between Physics and Number Theory: Convergent Theses from Ostrowski's Theorem to Adelic Quantum Mechanics. [11] https://papers.qnfo.org/papers/all-things-are-one. QNFO: All Things Are One.