

On the Impossibility of Deriving Quantitative Intra-Multiplet Splittings from Symmetry-Complete Dynamics Alone

A General No-Go Theorem

(c)July 2026 Sergej Materov*

Abstract

During repeated attempts to derive the observed three-generation flavor hierarchy solely from an exact graph symmetry, every construction either preserved exact multiplet degeneracy or required external symmetry-breaking input. This suggested that the obstruction was intrinsic to symmetry-complete theories rather than to any particular graph construction, motivating the general no-go theorem proved here.

A recurring claim across several research programs - discrete flavor-symmetry model building (A_n , S_n , $\Delta(27)$ family symmetries), Froggatt–Nielsen-type charge assignments, and combinatorial or graph-based approaches to emergent spacetime and particle content - is that the observed hierarchy of fermion masses, mixing angles, or analogous multiplet splitting's can be derived from symmetry considerations alone, without an independently motivated symmetry-breaking sector. We show this is impossible in a precise, model-independent sense. If a theory is defined entirely by a symmetry group S , a covariant representation Γ of S , and dynamics H_{gen} covariant under S , then no numerical value can be derived within that theory for the splitting between components of a non-trivial irreducible multiplet of observables, beyond the trivial statement that the splitting is exactly zero when the theory's distinguished state is S -invariant. If the distinguished state instead spontaneously breaks S , the direction of breaking is undetermined by the theory, and even granting a direction, the typical magnitude of splitting predicted by a symmetry-neutral (Haar-random) prior is $O(1)$ between multiplet components — so any large observed hierarchy is itself evidence of a further, non-symmetry input. The result is an exact operator-level sharpening of Curie's symmetry principle (1894) via the Wigner–Eckart theorem (1927), together with the standard vacuum-selection problem of spontaneous symmetry breaking (Goldstone 1961; Goldstone, Salam & Weinberg 1962). It applies to any “symmetry-complete” theory - continuum or discrete, field-theoretic or lattice-based and identifies exactly which additional, non-symmetry ingredient any such theory must supply.

1. Historical Context

The results below are not new physics. They are an exact, operator-level version of two classical results, stated together with a precision that is occasionally lost in informal usage.

Curie's principle (1894): “When certain causes produce certain effects, the symmetry elements of the causes must be found in the effects.” Equivalently: an effect cannot have less symmetry than its cause. The theorem below gives the precise, quantitative operator statement of this principle for matrix elements, and identifies exactly what must be added to a symmetric cause to obtain an asymmetric effect.

The Wigner–Eckart theorem (Wigner 1927; Eckart 1930): a matrix element of a spherical-tensor (irreducible-representation) operator between angular-momentum eigenstates factorizes into a Clebsch–Gordan coefficient and a reduced matrix element; the coefficient vanishes when the representation content does not permit the transition. Section 2 below is the same statement stripped of the specific group $SO(3)$ and specific coordinates, applied to the case of a diagonal matrix element in a single S -invariant state.

*Independed Researcher. ORCID: 0009-0001-3398-9906. E-mail: sergejmaterov2@gmail.com

Spontaneous symmetry breaking and vacuum degeneracy (Goldstone 1961; Goldstone–Salam–Weinberg 1962): when the ground state does not share the symmetry of the dynamics, the vacuum manifold forms an orbit of the symmetry group, and no rule internal to the dynamics selects a point on that orbit. Section 3 restates this as a precise obstruction for any theory that admits no external input.

What is new here is only the packaging: a single, short, model-independent theorem, together with an explicit typicality bound (Section 4) quantifying how large a splitting a generic, non-fine-tuned direction of breaking can plausibly produce. This is intended as a citable reference result to be invoked whenever a symmetry-complete theory is asked to explain a specific, large numerical hierarchy.

1.1 Scope of Applicability

The theorem is stated abstractly (§2) so as to apply uniformly to several otherwise unrelated research programs that share the same logical structure — a claim that a symmetry group, together with covariant dynamics, suffices to fix a quantitative splitting pattern:

- Discrete flavor-symmetry model building (A_4 , S_3 , S_4 , $\Delta(27)$, and related family symmetries), where fermion mass hierarchies and mixing angles are sought from the representation content of a flavor group (Altarelli–Feruglio and the subsequent literature). Most serious constructions in this line already introduce an explicit flavon sector with its own potential, i.e. they already supply the external input this theorem shows to be necessary, typically without stating the requirement in this general form.
- Froggatt–Nielsen-type constructions, where a discrete charge assignment is introduced by hand specifically to generate hierarchies - an explicit acknowledgment, in practice, of Theorem B/Proposition C below.
- Combinatorial and graph-based approaches to emergent spacetime and matter content, including loop quantum gravity and spin-network models (Rovelli, Smolin), causal set theory, and related programs where macroscopic symmetries (isotropy, gauge structure, particle multiplets) are proposed to arise purely from the superposition or coarse-graining of a symmetric microscopic substrate - an application, in spirit, of Curie's principle at the Planck scale. The theorem applies whenever such a program further claims to derive quantitative (not merely qualitative) splittings within an emergent multiplet from the microscopic symmetry alone.
- Grand unified models that attempt to fix Yukawa hierarchies from the group-theoretic content of a single unifying representation, without an independently motivated Higgs/flavon sector beyond the minimal one required for gauge symmetry breaking.

The Standard Model itself is not in scope: its Yukawa couplings are free parameters of the Lagrangian, not quantities claimed to be derived from the gauge group's representation theory, so it is not “symmetry-complete” in the sense of Definition 1. Likewise, any theory that already, explicitly, introduces a second ingredient beyond $(S, \Gamma, \mathcal{A}, \hat{H}_{gen})$ is a flavon field, an extra-dimensional profile, an explicit texture - is not constrained by this theorem beyond the (now discharged) requirement that such an ingredient be present.

2. Setup

Let:

- S - a group (finite, or compact Lie - the proof is identical; for Lie groups replace group-averages below by integrals against the Haar measure).
- $\Gamma: S \rightarrow U(\mathcal{H})$ - a unitary representation of S on the theory's Hilbert space $\mathcal{H}$.
- $\mathcal{A}$ - an algebra of observables closed under S -conjugation: for all $\hat{O} \in \mathcal{A}$ and $g \in S$, $\Gamma_g \hat{O} \Gamma_g^{-1} \in \mathcal{A}$.
- H_{gen} - a generator of dynamics, S -covariant: $\Gamma_g \hat{H}_{gen} \Gamma_g^{-1} = \hat{H}_{gen}$ for all $g \in S$.
- Ω - the state singled out by the theory through an S -covariant rule and no other input (e.g. the unique ground state of $\hat{H}_{gen}$, or the unique S -invariant Gibbs/stationary state) - not a state chosen by any additional, non- S -derived criterion.

Definition 1 (Symmetry-complete theory). The tuple $\mathcal{T} = (S, \Gamma, \mathcal{A}, \hat{H}_{gen}, \Omega)$ is called symmetry-complete if every ingredient above is determined by S and by nothing else - in particular, Ω is not selected by any datum external to $(S, \Gamma, \mathcal{A}, \hat{H}_{gen})$.

Definition 2 (Non-trivial multiplet). A finite set of observables $\{\hat{O}_a\}_{a=1}^d \subseteq \mathcal{A}$ is a d -dimensional multiplet transforming under an irreducible representation ρ of S if $\Gamma_g \hat{O}_a \Gamma_g^{-1} = \sum_i D(g)_{ia} \hat{O}_i$ for all $g \in S$, where D is the matrix form of ρ . The multiplet is non-trivial if ρ is not the trivial (identity) representation.

3. Theorem A: Exact Vanishing (S-invariant Ω)

Theorem A.

Let $\mathcal{T} = (S, \Gamma, \mathcal{A}, \hat{H}_{gen}, \Omega)$ be symmetry-complete, with $\Gamma_g \Omega = \Omega$ for all $g \in S$ (automatic when Ω is a non-degenerate ground state of the covariant H_{gen} , since Γ_g then maps the ground eigenspace to itself and non-degeneracy forces $\Gamma_g \Omega = e^{i\theta(g)} \Omega$, and an overall phase can be absorbed by redefining Γ). Let $\{\hat{O}_a\}$ be any non-trivial multiplet (Definition 2). Then

$\langle \Omega | \hat{O}_a | \Omega \rangle = 0$ for every $a = 1, \dots, d$.

The value is exactly zero, not merely small: no refinement of the dynamics $\hat{H}_{gen}$ within $\mathcal{A}$ can change it, because the vanishing follows from representation theory alone.

Proof.

Define $v_a := \langle \Omega | \hat{O}_a | \Omega \rangle$. Using $\Gamma_g \Omega = \Omega$ and unitarity of Γ_g :

$$v_a = \langle \Omega | \hat{O}_a | \Omega \rangle = \langle \Gamma_g \Omega | \hat{O}_a | \Gamma_g \Omega \rangle = \langle \Omega | \Gamma_g^{-1} \hat{O}_a \Gamma_g | \Omega \rangle = \sum_i D(g)_{ia} \langle \Omega | \hat{O}_i | \Omega \rangle = \sum_i D(g)_{ia} v_i, \quad \forall g \in S.$$

Hence the vector $v = (v_1, \dots, v_d)$ is invariant under the action of S in representation ρ : $v = \rho(g)v$ for all g . Averaging over the group,

$$v = \left(\frac{1}{|S|} \sum_{g \in S} \rho(g) \right) v \quad [\text{or } v = \int_S \rho(g) d\mu(g) \cdot v, \text{ for compact Lie } S].$$

The averaged operator on the right is exactly the projector onto the trivial isotypic component of ρ . By Schur orthogonality, this projector is identically zero whenever ρ contains no trivial subrepresentation - which holds by hypothesis, since ρ is non-trivial and irreducible. Therefore $v = 0$. ■

This is precisely the Wigner–Eckart theorem specialized to a diagonal matrix element in a single S -invariant state: the “Clebsch–Gordan coefficient” governing $\text{trivial} \otimes \rho \otimes \text{trivial} \rightarrow \text{trivial}$ vanishes unless $\rho = \text{trivial}$.

Corollary (exact Curie principle). No amount of computation performed strictly within $\mathcal{A}$ and H_{gen} can produce a non-zero splitting among the components of a non-trivial multiplet, when the theory's state is S -invariant. The vanishing is a logical consequence of the representation content, not a numerically small or “fine-tuned-away” quantity - increasing precision, computing power, or model detail cannot move it off zero.

Remark A.1 (the UV-completion regress does not evade the theorem). A natural objection is that potential parameters (e.g. a Higgs-sector m^2, λ) need not be “external” in the informal sense if they are themselves the output of some higher-energy dynamics. This does not evade Theorem A: either that UV dynamics is itself symmetry-complete under some (possibly larger) group S' , in which case the theorem applies verbatim one level up - the obstruction simply relocates, it does not disappear - or the UV theory supplies its parameters as free/measured inputs, in which case it is precisely the external datum the theorem requires, merely pushed to a higher scale. There is no finite regress that terminates in a symmetry-complete theory with a non-trivial multiplet splitting; at whatever level the regress stops, that level must contain a non- S -derived input.

Remark A.2 (hidden external input via dynamical scale generation). External input need not take the form of an explicit dimensionful parameter written into H_{gen} . Dynamical scale generation - e.g. dimensional transmutation in an asymptotically free gauge theory, where a scale Λ is generated from the running of a classically dimensionless coupling $g(\mu)$ - also requires external input, in the form of a boundary condition $g(\mu_0)$ fixing the integration constant of the renormalization-group flow. Nothing in $(S, \Gamma, \mathcal{A}, \hat{H}_{gen})$ alone selects this boundary value; it is again a datum $s \notin (S, \Gamma, \mathcal{A}, \hat{H}_{gen})$, simply supplied at the level of a renormalization scheme rather than a Lagrangian mass term. The theorem is agnostic to the form the external input takes is explicit parameter, flavon vacuum expectation value, or RG boundary condition - only to its necessity.

4. Theorem B: The Vacuum-Selection Obstruction (degenerate Ω)

Theorem A assumed $\Gamma_g \Omega = \Omega$. Suppose instead Ω is degenerate and the degenerate manifold itself transforms non-trivially under S (spontaneous symmetry breaking to a subgroup $H \subset S$). Theorem A no longer applies directly, and a specific representative Ω_0 of the degenerate manifold can have non-zero $v_a = \langle \Omega_0 | \hat{O}_a | \Omega_0 \rangle$.

Theorem B.

Under the hypotheses above, the set of states compatible with $\mathcal{T}$ is the full orbit $\{\Gamma_g \Omega_0 : g \in S\} \cong S/H$. The data $(S, \Gamma, \mathcal{A}, \hat{H}_{gen})$ contain no rule selecting a particular point on this orbit: for any two $g_1, g_2 \in S$, the states $\Gamma_{g_1} \Omega_0$ and $\Gamma_{g_2} \Omega_0$ give physically isomorphic (S -conjugate) theories. Selecting one particular Ω_0 therefore requires a datum $s \notin (S, \Gamma, \mathcal{A}, \hat{H}_{gen})$.

Justification.

This is the standard vacuum-selection problem of spontaneous symmetry breaking (Goldstone 1961; Goldstone–Salam–Weinberg 1962): a symmetric Hamiltonian with a degenerate, symmetry-related family of ground states does not itself specify which ground state is realized. In the thermodynamic / continuum limit, physical selection occurs via an infinitesimal explicit symmetry-breaking seed, taken to zero only after the limit and the direction of that seed is external data by construction, not a consequence of $\hat{H}_{gen}$. Equivalently: distinct points of the orbit lie in different superselection sectors, related by the unbroken part of the symmetry, and no operator in $\mathcal{A}$ connects them or prefers one over another.

Remark B.1 (Theorem B is strictly an infinite-volume statement; finite systems default to Theorem A). For finite S acting on a finite-dimensional $\mathcal{L}$, exact ground-state degeneracy across a full non-trivial orbit is non-generic: level splitting between classically degenerate configurations typically survives as an exponentially small but non-zero tunneling amplitude (the finite-size transverse-field Ising chain is the textbook example its true ground state remains unique and symmetric at any finite size, with the symmetry-broken doublet only becoming exactly degenerate as the system size, and hence the tunneling suppression, is taken to infinity). Consequently, for finite systems Theorem A - not Theorem B - is generically the operative statement: Ω is typically unique and S -invariant, and the exact-zero conclusion of Theorem A applies literally, not merely as an idealization. Theorem B becomes physically operative only in the strict thermodynamic or continuum limit, where tunneling between orbit points is suppressed to exactly zero rather than merely exponentially small. A theory formulated on a finite substrate that reaches its symmetry-breaking claims only after an idealized infinite-volume limit should state this limit explicitly, since the finite-volume theory it is built from is, before that limit, governed by the stronger and more restrictive Theorem A.

5. Proposition C. Strengthened: Exact Typicality Bound on the Magnitude of Breaking

Theorem B shows that the direction of symmetry breaking is external input. Proposition C, in its old form, argued heuristically - via 20,000 Monte Carlo samples for the case $d = 2$ - that even a typical (non-fine-tuned) magnitude of breaking predicts only an $O(1)$ splitting ratio between multiplet components, so that a large observed hierarchy is itself evidence of an atypical, further input. Below, this claim is upgraded from a numerical illustration to an exact theorem for $d = 2$, together with a precision-quantified extension to general d , and an explicit statement of the maximum-entropy justification for the prior itself.

5.1 The prior is the unique maximum-entropy choice, not an arbitrary convention

Let the symmetry-breaking perturbation acting within a d -dimensional real multiplet be modeled by a real symmetric matrix M on the representation space. “Symmetry-neutral” is made precise as follows:

Lemma (maximum entropy). Among all probability distributions on real symmetric $d \times d$ matrices that (i) are invariant under conjugation by every element of $O(d)$ - i.e. carry no preferred direction in the multiplet space — and (ii) have a fixed value of $E[\text{Tr}(M^2)]$ (a fixed overall breaking scale), the distribution of maximum differential entropy is the Gaussian Orthogonal Ensemble,

$$P(M) \propto \exp\left(-\frac{\text{Tr}(M^2)}{2\sigma^2}\right).$$

Proof. Among all distributions on $\mathbb{R}^{\frac{d(d+1)}{2}}$ (the space of symmetric matrices, in coordinates given by the independent entries) with fixed covariance, the entropy-maximizing distribution is Gaussian (standard maximum-entropy result for a quadratic constraint). Imposing $O(d)$ -invariance is a linear constraint on the family of covariance matrices consistent with (i); the unique covariance compatible with $O(d)$ -invariance is proportional to the identity on the appropriately weighted entry space, which is exactly the GOE normalization $\text{Var}(M_{ii}) = \sigma^2$, $\text{Var}(M_{ij}) = \frac{\sigma^2}{2}$ ($i < j$). ■

This closes the gap in the old Proposition C, where the choice of a Haar/GOE prior was asserted rather than derived: it is not one convention among several, but the unique rotation-invariant, maximum-entropy prior at fixed breaking scale - precisely the correct formalization of “no preferred direction, no further structure assumed.”

5.2 Exact result for $d = 2$

Let $M = [[x, z], [z, y]]$ with $x, y \sim N(0, \sigma^2)$ and $z \sim N(0, \frac{\sigma^2}{2})$, independent (the GOE(2) normalization fixed by §5.1). Write the eigenvalues as $\lambda_{1,2} = m \pm r$, where $m = \frac{x+y}{2}$ and $r = \sqrt{u^2 + z^2}$ with $u = \frac{x-y}{2}$. A direct change of variables gives, independently of σ ,

$$m \sim N\left(0, \frac{\sigma^2}{2}\right), \quad r \sim \text{Rayleigh}\left(\frac{\sigma^2}{2}\right)$$

[density $f(r) = \frac{2r}{\sigma^2} \exp\left(-\frac{r^2}{\sigma^2}\right)$, $r \geq 0$], and m, r are independent. Two standard consequences:

(a) Splitting distribution is exactly Wigner’s surmise. The splitting $s = \lambda_1 - \lambda_2 = 2r$ has density $f(s) = \frac{s}{2\sigma^2} \exp\left(-\frac{s^2}{4\sigma^2}\right)$; after rescaling to unit mean spacing this is exactly the Wigner surmise $P(s) = \frac{\pi s}{2} \exp\left(-\frac{\pi s^2}{4}\right)$ - not an approximation, as it is for GOE(d) with $d > 2$, but an exact identity for $d = 2$. ($E[s] = \sigma\sqrt{\pi}$, confirmed both analytically and by 20-million-sample simulation to five significant figures.)

(b) Exact closed form for the splitting ratio. Define $R = \max(|\lambda_1|, |\lambda_2|) / \min(|\lambda_1|, |\lambda_2|)$, the quantity of physical interest (e.g. the ratio of two mass-squared eigenvalues within a doublet). Setting $v = r/|m|$, one finds that R is a function of v alone $R(v) = \frac{1+v}{|1-v|}$ and, by a ratio-distribution computation using the two densities above,

$$f_v(v) = v(I+v^2)^{-3/2}, \quad v \geq 0, \quad F_v(v) = 1 - (I+v^2)^{-1/2}.$$

Integrating over the two branches of $R(v)$ gives the exact, closed-form tail probability

$$P(R > r_0) = \sqrt{\frac{2}{1+r_0^2}}, \quad r_0 \geq 1.$$

Consequently, exactly:

$$\text{median}(R) = \sqrt{7} \approx 2.646, \quad 90\text{th pctile}(R) = \sqrt{199} \approx 14.107, \quad 99\text{th pctile}(R) = \sqrt{19999} \approx 141.42,$$

$$P(R > 20) = \sqrt{\frac{2}{401}} \approx 7.062\%, \quad P(R > 200) = \sqrt{\frac{2}{40001}} \approx 0.707\%.$$

These values were cross-checked by 20-million-sample Monte Carlo (median 2.6444, 90th 14.109, 99th 141.97, $P(R>20)$ 7.066%, $P(R>200)$ 0.711%) and by independent deterministic numerical quadrature of the joint density (matching the closed form to 5+ significant figures at every tested r_0). All numbers in the original Proposition C

(median ≈ 2.4 , 90th ≈ 12.6 , 99th ≈ 126 , from 20,000 samples) are now superseded by the exact values above; the qualitative claim survives, and is now a theorem rather than a simulation (see the GitHub repository in the "Reference" section).

5.3 General d: quantified numerics, and an honest correction to the original claim

For $d > 2$ no equally simple closed form is available (the joint eigenvalue density of $\text{GOE}(d)$ is known exactly - the Weyl/Wigner formula - but the distribution of the extremal ratio $\frac{\max|\lambda_i|}{\min|\lambda_i|}$ does not reduce to an elementary integral). High-precision Monte Carlo (3×10^6 independent $\text{GOE}(d)$ draws per row, generated with the exact normalization of §5.1) gives:

d	median R	90th pct.	99th pct.	P(R>20)	P(R>200)
2 (exact)	2.646	14.107	141.42	7.06%	0.707%
3	5.01	26.97	269.8	13.46%	1.35%
4	7.22	38.99	390.3	19.32%	1.95%
5	9.57	51.70	517.5	25.42%	2.57%
6	11.83	63.98	644.4	31.12%	3.22%

Table 1. Splitting-ratio typicality by multiplet dimension d (GOE , maximum-entropy prior).

This is an important refinement the original Proposition C did not make: the typical splitting ratio grows systematically with the dimension of the multiplet. This is an expected extreme-value-statistics effect $\frac{\max|\lambda_i|}{\min|\lambda_i|}$ is an extremal-order statistic over d eigenvalues, and such statistics grow with sample size even under a fixed, non-fine-tuned generating distribution - and it means the “ $O(1)$ is typical” language of the original text is only strictly accurate at $d = 2$. For a 3-generation flavor multiplet ($d = 3$, the case of direct interest to the paper’s stated motivation), the correct typicality benchmark is median ≈ 5 , not $O(1)$; a 99th-percentile symmetry-neutral ratio for $d = 3$ is already ≈ 270 .

For asymptotically large d (many-generation or continuum multiplets), the appropriate object is not the global extremal ratio but the local ratio of consecutive level spacings, whose distribution is known in closed form via the Atas–Bogomolny–Giraud–Roux surmise (Atas et al., Phys. Rev. Lett. 110, 084101, 2013):

$$P\beta(\rho) = \frac{1}{Z\beta} \cdot \frac{(\rho + \rho^2)^\beta}{(1 + \rho + \rho^2)^{1+\frac{3\beta}{2}}}$$

with $\beta = 1$ for GOE , $\rho = \frac{\min(s_n, s_{n+1})}{\max(s_n, s_{n+1})} \in [0, 1]$ the ratio of consecutive spacings s_n . This is the standard tool for characterizing typical level statistics in large random-matrix multiplets and should be cited as the correct asymptotic completion of the $d = 2$ result above, rather than re-deriving a new bulk-statistics result from scratch.

Numerical verification. The $d = 3$ median of Table 1 was independently confirmed by deterministic Gauss–Hermite quadrature over the exact $\text{GOE}(3)$ joint eigenvalue density (90^3 -node tensor grid), giving median = 5.0137 against the Monte Carlo value 5.01 (Table 1); tail probabilities, however, require adaptive resolution near the near-degenerate region $\min|\lambda_i| \rightarrow 0$ and were therefore obtained by Monte Carlo, as in Table 1. Separately, the bulk consecutive-spacing ratio was simulated directly for $d = 30, 60, 100$ and found to converge to $\langle \rho \rangle \approx 0.530\text{--}0.532$, consistent with the GOE ($\beta = 1$) value $\langle r \rangle \approx 0.5307$ of Atas et al. (2013), confirming that Table 1 (finite, small d)

and the asymptotic surmise cited above (large d) describe the same underlying random-matrix ensemble across its full range of applicability.

5.4 Proposition C

Proposition C (revised). Let ρ be a d -dimensional ($d \geq 2$) real irreducible representation of S , and let the symmetry-breaking perturbation within the multiplet be modeled by a matrix drawn from the unique $O(d)$ -invariant, maximum-entropy prior at fixed breaking scale (§5.1), i.e. the Gaussian Orthogonal Ensemble. For $d = 2$, the tail probability of the resulting eigenvalue-splitting ratio R is exactly $P(R > r_0) = \sqrt{\frac{2}{1+r_0^2}}$; in particular the median ratio is exactly $\sqrt{7} \approx 2.65$, and ratios exceeding 200 occur with prior probability exactly $\sqrt{\frac{2}{40001}} \approx 0.71\%$. For $d > 2$, the typical ratio grows systematically with d (Table 1, §5.3); the qualitative conclusion - that ratios far in excess of the d -appropriate typical scale constitute evidence of an atypical, non-symmetry-derived magnitude - holds for every d , but the quantitative threshold must be read off the d -dependent distribution rather than treated as a universal $O(1)/O(10)$ benchmark.

6. Main Result

Theorem (No-Go for Quantitative Multiplet Splitting).

Let $\mathcal{T} = (S, \Gamma, \mathcal{A}, \hat{H}_{gen}, \Omega)$ be symmetry-complete (Definition 1), and let $\{\hat{O}_a\}$ be any non-trivial irreducible multiplet of observables (Definition 2). Then:

- (i) If Ω is S -invariant: all components of the multiplet are exactly degenerate, $\langle \Omega | \hat{O}_a | \Omega \rangle = 0$ identically (Theorem A) - not approximately, not “negligibly small,” but a forced logical consequence of the representation content.
- (ii) If Ω spontaneously breaks S : the direction of breaking is undetermined by $\mathcal{T}$ and requires an external datum (Theorem B).
- (iii) Even granting a direction externally: the magnitude required to reproduce a large observed hierarchy is itself, generically, atypical relative to a symmetry-neutral prior, and so constitutes further information beyond the bare symmetry-breaking mechanism (Proposition C).

Hence: quantitative prediction of intra-multiplet splittings requires at minimum one (non-degenerate Ω broken by hand - not permitted within $\mathcal{T}$ at all, cf. (i)) or, more precisely, at least two independent empirical inputs beyond $(S, \Gamma, \mathcal{A}, \hat{H}_{gen})$ in the spontaneously-broken case: a direction and a magnitude. No symmetry-complete theory, however detailed its covariant dynamics, can supply these from S alone.

7. Scope: What the Theorem Does Not Forbid

The result is deliberately narrow, and it is important to state its limits precisely so that it is not mis-cited as a broader impossibility result than it is.

- (a) Qualitative representation-theoretic structure is unaffected. The decomposition of a reducible representation into irreducibles (e.g. which observables group into which multiplets, and their dimensions) is a fact about S

alone and is fully derivable within $\mathcal{T}$; the theorem constrains only the numerical splitting within an already-identified non-trivial multiplet.

(b) Observables that are themselves S-invariant (transform in the trivial representation) are entirely unconstrained by this theorem - their expectation values in Ω are generically non-zero and may be computed from $\hat{H}_{gen}$ without any external input.

(c) The theorem does not address whether a symmetry-breaking mechanism exists (i.e. whether Ω can be made to spontaneously break S at all, as a function of some control parameter in $\hat{H}_{gen}$) - only what happens once such a mechanism is in place: existence of a phase transition is a separate, model-dependent dynamical question outside this theorem's scope.

(d) The theorem says nothing about multiplets built from a different symmetry $S' \neq S$ under which Ω is not distinguished as invariant/covariant in the relevant sense - e.g. a theory may derive one sector's structure (say, a gauge group or a spacetime dimensionality) fully from S while a separate, non-trivial-under-S multiplet (say, a flavor hierarchy) remains subject to Theorem A/B. The two are logically independent, and a theorem-A obstruction for one multiplet is not evidence against the rest of a symmetry-complete construction.

8. A Related but Independent Obstruction: Algorithmic Undecidability

Theorems A and B obstruct quantitative multiplet splittings by an exact representation-theoretic vanishing (or, in the broken case, an unresolved orbit-selection problem). A second, logically independent obstruction to inferring macroscopic quantitative structure from microscopic data exists in a different setting and by a completely different mechanism: Cubitt, Pérez-García & Wolf (2015) proved that the spectral-gap problem - whether a translationally invariant, local Hamiltonian on an infinite lattice is gapped or gapless in the thermodynamic limit - is algorithmically undecidable in general, by an explicit encoding of Turing-machine halting into a two-dimensional quantum many-body model built from Wang/Robinson tilings. This shows that even complete, exact knowledge of a local, translationally symmetric microscopic Hamiltonian does not, in general, computably determine a basic macroscopic spectral property.

We record this as a separate, independent result and do not fold it into Theorems A–B: undecidability is a statement about computability (there exists no algorithm that decides the property for all inputs in the relevant class), whereas Theorems A–B are statements about exact vanishing and orbit-selection that hold instance-by-instance and require no appeal to computability theory. A theory could in principle evade Cubitt–Pérez-García–Wolf undecidability (e.g. by lying outside the reduction's model class) while still being fully subject to Theorem A, and vice versa. The two results are complementary illustrations of the same broad moral - that microscopic completeness does not imply macroscopic derivability - arrived at by unrelated mathematics, and citing one is not a substitute for the other.

(As an aside, with no bearing on the argument above: S. Hawking's 2002 lecture “Gödel and the End of Physics” speculated, on general Kurt-Gödel-inspired grounds, that a complete physical theory may be unattainable in principle - a suggestive remark (in the spirit of popular science), not a mathematical result, and unrelated in mechanism to either Theorems A–B or the Cubitt–Pérez-García–Wolf reduction.)

9. Summary Statement

A theory built from symmetry and covariant dynamics alone can tell you the qualitative shape of a multiplet - how many components, what representation - but it cannot, even in principle, tell you the numbers by which those components differ; that information is either exactly zero (unbroken case) or requires at least a direction and a scale supplied from outside the theory (broken case).

References

- Curie, P. (1894). “Sur la symétrie dans les phénomènes physiques, symétrie d'un champ électrique et d'un champ magnétique.” *Journal de Physique Théorique et Appliquée*, 3(1), 393–415.
- Wigner, E. (1927). “Einige Folgerungen aus der Schrödingerschen Theorie für die Termstrukturen.” *Zeitschrift für Physik*, 43, 624–652.
- Eckart, C. (1930). “The Application of Group Theory to the Quantum Dynamics of Monatomic Systems.” *Reviews of Modern Physics*, 2, 305–380.
- Goldstone, J. (1961). “Field Theories with ‘Superconductor’ Solutions.” *Il Nuovo Cimento*, 19(1), 154–164.
- Goldstone, J., Salam, A., & Weinberg, S. (1962). “Broken Symmetries.” *Physical Review*, 127(3), 965–970.
- Serre, J.-P. (1977). *Linear Representations of Finite Groups*. Springer GTM 42. (Standard reference for Schur orthogonality and projector formulas used in Section 3.)
- Cubitt, T. S., Pérez-García, D., & Wolf, M. M. (2015). “Undecidability of the Spectral Gap.” *Nature*, 528, 207–211.
- Altarelli, G., & Feruglio, F. (2010). “Discrete Flavor Symmetries and Models of Neutrino Mixing.” *Reviews of Modern Physics*, 82, 2701–2729.
- Froggatt, C. D., & Nielsen, H. B. (1979). “Hierarchy of Quark Masses, Cabibbo Angles and CP Violation.” *Nuclear Physics B*, 147(3–4), 277–298.
- Hawking, S. W. (2002). “Gödel and the End of Physics.” Dirac Centennial Lecture, Cambridge (informal remarks; not a peer-reviewed result - cited here only as historical context, §8).
- Materov, S. (2026). GitHub repository: “Computational Appendix - Proposition C (GOE Splitting-Ratio Typicality)”, <https://github.com/SergejMaterov/SymmetryNoGo>