Deriving the Particle Zoo from Observer Consistency

B. Müller Alexander Osika Kai Xue Mario Poneder

June 8, 2026

Paper release: `r1465` **Released:** June 8, 2026

Abstract

Observer-Patch Holography turns the particle problem into a forward calculation from one local screen-geometry fixed point, not from a fitted particle-sector parameter. The screen is a regulator and symmetry chart whose finite geometry organizes the A_5- and E_8-type data. The local pixel closure proceeds through five physics layers: unification scale, heat-kernel pixel closure, running transport, electroweak mixing, and Ward-projected Thomson transport. The outer/inner pixel equation leaves no local tuning freedom for the fine-structure constant:

$$P_\star = \varphi + \frac{\sqrt{\pi}}{A_T(P_\star)}.$$

The fixed point gives $\alpha^{-1}(0) = 137.035999177(21)$ and $P \simeq 1.6309682094$. The same pixel scale propagates into structural massless carriers, the weak sector, the Higgs/top row, selected-class running quarks, and a weighted-cycle neutrino absolute-mass branch. Hadrons require separate treatment: source-only strong-binding calculations are kept distinct from empirical $e^+e^- \to$ hadrons closure checks. The status table separates validation rows, charged-lepton witnesses, source-only checks, and empirical hadron closure checks.

1 Introduction

Observer-Patch Holography asks a concrete follow-up question after the gauge structure is fixed. Can one local screen constant organize the particle spectrum, or are the masses and mixings a collection of unrelated inputs?

The structural companion papers start from observer consistency on a finite screen. Physical data are organized in local patch algebras, nearby observers must agree on overlaps, realized states obey a local maximum-entropy and refinement principle, generalized entropy is recoverable, and the realized low-energy branch is fixed by Minimal Admissible Realization (MAR). On the certified support-visible BW/geometric branch, that basis yields the Lorentzian and Einstein branches; on the realized compact-gauge branch it yields the Standard Model gauge structure

$$\frac{\mathrm{SU}(3)\times \mathrm{SU}(2)\times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_c = 3, \qquad N_g = 3.$$

The Lorentzian branch, Einstein recovery, and the realized Standard Model gauge structure written above are imported from [4] with their stated branch conditions. That same paper also derives the hypercharge lattice and the realized color triplet $N_c = 3$ and generation count $N_g = 3$. This paper uses those results as input and summarizes only the downstream particle-spectrum continuation. If that derivation is correct, it should also constrain the particle spectrum.

The particle branch studied here is a one-fixed-point forward reconstruction problem with an explicit closure matrix. Instead of treating masses, mixings, and Yukawa parameters as independent inputs, it tries to propagate one dimensionless pixel ratio P through the spectrum. The pipeline claims closure only for the structural massless carriers, the Higgs/top row on its declared quantitative surface, the selected-class quark theorem, and the weighted-cycle neutrino absolute-attachment branch. The P-closure root, the electroweak W/Z rows, charged leptons, and hadrons have the sector-specific boundary statuses recorded below. This paper develops the derivation chain from the axioms and candidate P trunk to the reported particle outputs, with the closure boundary stated by sector.

1.1 Why this particle content is selected

On the declared branch, particles are not inserted by hand as a list of elementary ingredients. They are the stable observer-visible excitations and carrier modes that are left once three layers of structure have been fixed:

1. overlap consistency on the observer-patch network;
2. compact gauge reconstruction from the theorem-produced transportability criterion, fixed-cutoff bosonic sector category, refinement/fiber ladder, and realized MAR-admissible witness data;
3. minimal admissible realization of the low-energy branch.

The first layer says what kinds of local data can be compared consistently across neighboring patches. The second assembles persistent zero-obstruction sector data into a compact gauge structure through the compact paper's transportability, category, and refinement/fiber theorems. The third selects which of the admissible low-energy branches is physically realized. At that point one is not choosing an arbitrary particle catalog. One is reading off the carrier and matter content of the realized branch.

Recovering Relativity and Standard Model Structure from Observer Overlap Consistency [4] is therefore central to this draft. That paper does the structural work that fixes the realized low-energy package

$$\frac{\mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)}{\mathbb{Z}_6},\qquad N_c=3,\qquad N_g=3,$$

together with one admissible Higgs doublet and the realized chiral matter content. The construction does not begin by guessing photons, gluons, quarks, leptons, and weak bosons and then asking how to fit them. It first derives the realized gauge-and-matter branch, and that branch determines which elementary carriers and matter families are available at low energy.

This also explains the phrase "and no others." The structural OPH claim concerns the realized low-energy branch selected by the stated admissibility conditions. Once one imposes those conditions, the realized low-energy branch is the Standard-Model branch. Additional light gauge bosons, additional light Higgs multiplets, extra light chiral families, or low-energy supersymmetric partners do not appear on that realized branch because they would enlarge the admissible package that MAR selects against. In that sense the observed light particle zoo is not an arbitrary menu; it is the minimal realized one.

The individual families then emerge for different reasons. The photon, gluons, and graviton are structural carriers: once the realized electromagnetic, color, and dynamical-metric branches are present, those massless carriers follow as symmetry-protected zeros. The weak bosons arise when the electroweak gauge sector is propagated through its quantitative closure branch. Quarks and leptons arise from the realized chiral matter package together with the three-generation structure $N_g = 3$. Their family splittings are then read from the deeper overlap-transport and excitation machinery, not from a second arbitrary postulate that says "copy the family three times and assign masses by hand." Hadrons come one layer later: once color is realized and confined, stable hadrons are composite readout channels of the quark/gluon sector.

The role of the pixel constant P is also important to state clearly. Here, P does not decide whether the photon exists, whether there are three colors, or whether quarks come in three generations. Those structural facts belong to the gauge-and-matter branch fixed before the quantitative closure step. What P does is set the shared quantitative scale on which the realized particle content is read out numerically. So the logic of the paper is:

$$\text{observer consistency} \to \text{realized gauge/matter branch} \to \text{particle content}$$
$$\to \text{common scale } P \to \text{quantitative spectrum.}$$

That is how the derivation explains why a universe with the realized branch and the shared pixel scale exhibits this particle zoo.

1.2 Logical basis and reading rule

The particle derivation uses the same OPH basis as the SM/GR derivation in Ref. [4]. Its five axioms are the screen net, overlap consistency, local MaxEnt with refinement stability, recoverable generalized entropy, and Minimal Admissible Realization (MAR). On that basis, the SM/GR derivation in Ref. [4] supplies the structural chain used here: compact gauge reconstruction as a classification step, MAR selection of the realized Standard Model quotient, the hypercharge lattice, the realized color triplet $N_c = 3$, and the generation count $N_g = 3$. Refs. [2, 3] supply the patch-net, repair, measurement, and regulated screen language used later when flavor transport and observer-facing readout enter.

The quantitative particle side adds one local closure variable. The common pixel ratio P, fixed on the synthesis paper's outer/inner closure branch [1], feeds the forward electroweak map and its descendants. The fixed point gives

$$\alpha^{-1}(0) = 137.035999177(21), \qquad \alpha(0) \simeq 0.00729735256433, \qquad P \simeq 1.6309682094.$$

The computation has a fixed source order: golden-ratio entropy balance gives φ, boundary Gaussian normalization supplies the $\sqrt{\pi}$ width, a trial P feeds the source map through unification, running, and electroweak anchoring, and Ward-projected electromagnetic transport gives the Thomson endpoint used by the outer/inner pixel fixed point. The value is forced because each ingredient in that equation is fixed by the screen balance, normalization, source map, and Ward-projected endpoint; any different α gives a different P and fails the outer/inner equality for the same cell. The source-side audit trunk records

$$P = 1.630972095694329018179678925611918842701 69, \qquad \alpha_{\text{cand}}^{-1} = 136.99483516462164945794999458578719326 2$$

with the detailed endpoint table in Section 4. The separate cosmic record-closure target $N_{\text{CRC}} = F(N_{\text{CRC}})$, with count representation $N_\star = \text{MAR} \arg\max_N [\log |\Omega_N^{\text{sc}}| - N]$, belongs to the cosmological-capacity branch and only to legacy capacity-level neutrino side estimates. Structural carriers, quantitative outputs, exact sidecars, and continuation lanes are therefore downstream branches of one common local fixed point and theorem checklist. The public fine-structure display row is the declared Ward-projected Thomson endpoint of the OPH fixed-point equation. The source-only audit row and empirical hadron closure rows are separate status-table records. Section 4 records that ledger in full.

2 P-Closure and the Reverse-Engineering Strategy

Before the claim-tier table, the reverse-engineering claim should be stated plainly. In ordinary phenomenological use, the Standard Model does not derive the observed particle masses and mixings from one common microscopic constant. It is usually presented with 19 free parameters in the minimal massless-neutrino theory, and with at least 26 once neutrino masses and mixing are included, depending on neutrino-sector conventions. The OPH particle derivation compresses that situation into one universal pixel ratio P from the outer/inner closure relation described in the synthesis paper. The same P must drive all downstream bosonic, quark, lepton, neutrino, and hadronic branches.

One clarification matters. The broader zero-input formulation targets the global screen capacity by

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}),$$

where $F(N)$ is the active cosmic record capacity read back by observers inside the universe supplied with capacity N. The count-density representation is

$$N_\star = \mathrm{MAR}\arg\max_N \left[\log|\Omega_N^{\mathrm{sc}}| - N\right],$$

equivalently, in differentiable form, as the unique stable fixed point of $T_\eta(N) = N + \eta\, d(\log|\Omega_N^{\mathrm{sc}}| - N)/dN$ under the stated strict-concavity certificate. Informally, it is the unique global capacity where the universe reads back its own boundary without deficit or slack, parallel to the local $P_\star$ fixed point where outer screen detuning equals the inner Thomson observation scale. The observed branch reads this same capacity as $N_{\mathrm{scr}} \simeq 3.31 \times 10^{122}$. The particle-spectrum derivation studied here uses the local pixel fixed point $P_\star$ in place of a long particle-by-particle parameter list.

2.1 Why a single P matters

The structural theorems determine *what kind* of particle world is realized, but not *what absolute scale in GeV* that world occupies. The overlap, modular, gauge, anomaly, and admissibility arguments recover the realized Standard Model branch, the exact hypercharge pattern, three generations, three colors, and the symmetry-protected massless carriers. But those structural results do not by themselves tell us the numerical values of the W boson mass, the Higgs boson mass, or the up-quark mass.

The role of the pixel closure is to supply one common quantitative scale variable for the entire downstream spectrum. In the synthesis paper, the fine-structure lane asks for the nonzero detuning of a holographic screen cell such that the cell's outer geometric displacement from perfect self-similar equilibrium equals the electromagnetic observation scale emitted by the universe living on that same screen. The outer side of the closure is

$$P = \varphi + \alpha_{\mathrm{in}}(P)\sqrt{\pi}.$$

The first-principles computation is a five-step source chain. First, the golden-ratio entropy balance of the local screen cell supplies $\varphi = (1+\sqrt{5})/2$. Second, boundary maximum-entropy normalization fixes the width of the leading electromagnetic detuning to $\sqrt{\pi}$. Third, a trial P is sent through the source map

$$M_U(P) = E_P e^{-2\pi} P^{1/6}, \qquad E_{\rm cell}(P) = \frac{E_P}{\sqrt{P}},$$

followed by the heat-kernel closure

$$\bar{\ell}_{\rm SU(2)}(t_2(P)) + \bar{\ell}_{\rm SU(3)}(t_3(P)) = \frac{P}{4},$$

which selects $\alpha_U(P)$ and the running family $\alpha_i(m_Z; P)$. Fourth, electroweak mixing gives the source anchor

$$a_0(P) = \alpha_2^{-1}(m_Z; P) + \frac{5}{3}\alpha_1^{-1}(m_Z; P).$$

Fifth, Ward-projected $\mathrm{U}(1)_Q$ transport gives the Thomson endpoint

$$A_T(P) = T_Q(a_0(P), F_{\rm src}(P)) = \alpha_{\rm em}^{-1}(0; P),$$

and the cell closes when

$$P = \varphi + \frac{\sqrt{\pi}}{A_T(P)}.$$

Thus α is the observer-supporting electromagnetic width of the local pixel, and its value is fixed by the unique solution of this outer/inner consistency equation. The OPH fixed-point readout $\alpha^{-1}(0) = 137.035999177(21)$ gives

$$P \simeq 1.6309682094.$$

The same equation is the local ruler for the downstream particle rows. The five-layer OPH audit trunk emits the source-side diagnostic point

$$P = 1.630972095694329018179678925611918842701 69, \qquad \alpha_{\rm cand}^{-1} = 136.99483516462164945794999458578719326 20$$

The detailed table records how the diagnostic trunk, endpoint residual, and source-spectral payload fit together. A separate hardware note reports an optical-cavity check of the same fixed-point geometry; this is treated as corroborating engineering evidence. Once the fixed point is set, the particle question becomes whether all downstream readouts can be written as

$$X_j = G_j(P)$$

with no new sector-specific constants inserted by hand. This paper studies exactly that downstream map.

2.2 How the electroweak branch works

In the implementation used here, the quantitative electroweak branch is the first major readout from P. Its public target-free output family is

$$\{M_W^{\rm pole},\ M_Z^{\rm pole},\ \alpha_{\rm em}^{-1}(q^2),\ \sin^2\theta_W(q^2),\ v\},$$

emitted from one declared input P on the printed running/matching/threshold/scheme package. The construction is straightforward in spirit: start from P, build the electroweak running family, select the physical carrier point on that family, and read off the W boson and Z boson pair from that selected point. The Higgs/top critical stage inherits that same electroweak core and does not introduce a new free-input sector. On the declared Ward-projected transport branch, the low-energy electromagnetic row is represented by the Thomson endpoint

$$\alpha_{\mathrm{Th}}^{-1}(P) = \lim_{q^2 \to 0} \alpha_{\mathrm{em}}^{-1}(q^2; P)$$

of that same electromagnetic transport family. The forward logical order on this branch is:

$$P \longmapsto (M_U(P), E_{\mathrm{cell}}(P)) \longmapsto \alpha_U(P) \longmapsto (t_U(P), t_{\mathrm{tr}}(P)) \longmapsto (t_2(P), t_3(P), v(P)) \longmapsto \alpha_i(\mu_*; P).$$

The forward transmutation certificate records that the same basis reconstructs the unified diffusion parameter $t_U(P) = 4\pi^2 \alpha_U(P)$ and transmutation exponent $t_{\mathrm{tr}}(P) = 2\pi/((N_c + 1)\alpha_U(P))$ as the pixel-closure solve itself. Measured electroweak data appear only at the validation surface.

The electroweak branch is the place where the common fixed point $P_\star$ emits the weak-sector rows. The published W/Z values are

$$M_W = 80.377 \text{ GeV}, \qquad M_Z = 91.18797809193725 \text{ GeV}.$$

The selected-carrier chart and the freeze-once coherent repair surface are recorded in the status ledger below.

2.3 What the electroweak branch fixes

On the declared electroweak running/matching/threshold/scheme surface, the pixel variable P fixes the source basis $(\alpha_{2,m_Z}, \alpha_{Y,m_Z}, \eta_{\mathrm{source}}, v)$. The target-free source-only repair theorem fixes the mass-side pair (W, Z) together with the source-locked running-family anchor (a_0, s_0, v) on the declared D10 surface from that source basis fixed by the forward solve. The electromagnetic row is physically read only after Ward projection to the unbroken $\mathrm{U}(1)_Q$ channel; its low-energy value is the $q^2 \to 0$ endpoint of the same transport family. The derivation also includes two explicit validation surfaces beneath the public theorem output: the exact selected-carrier chart and the freeze-once coherent repair surface. The same tiered logic is useful later in the paper because some flavor continuations are work in progress, whereas the electroweak quantitative-closure lane has an explicit declared-surface contract. Its source payload, same-scheme remainder, and interval-certificate records are public bookkeeping artifacts for the declared bridge from P to the fine-structure endpoint. On the paper surface, the transmutation factor is $\beta_{\mathrm{EW}} = N_c + 1$; overloaded β-ratios appear only on compare-only diagnostic readouts and are not part of the theorem contract.

3 Results at a Glance

The particle derivation carries the local pixel scale into the directly comparable rows below. Detailed scope statements are collected in the ledger sections that follow.

Two points are worth stating explicitly. First, the theorem rows and the exact-hit rows are not the same object. The theorem table carries selected-class quark closure and a corpus-limited charged no-go boundary, while the exact-hit surface above it also contains exact same-family charged and quark witnesses, the restricted transport-frame quark exact chain, and exact diagnostic sidecars. Second, "the bosons are finished" is too blunt. The photon, gluons, and graviton are structural zeros; the W boson and Z boson are compare-only frozen-adapter values; and the Higgs boson sits on the declared Higgs/top critical surface.

This claim-tier boundary governs the rest of the paper. Whenever a later chapter becomes detailed about a work in progress family, it should be read through this table. The table records the constructive chain and keeps its claim tier explicit.

Chain	OPH output	Role
P-closure and fine structure	$\alpha^{-1}(0) = 137.035999177(21)$, $P \simeq 1.6309682094$	local fixed-point scale
Structural massless carriers	$m_\gamma = m_g = m_{\text{grav}} = 0$	symmetry-protected zeros
Electroweak W/Z	80.377 GeV, 91.18797809193725 GeV	weak-sector validation pair
Higgs/top surface	$m_H = 125.1995304097179$ GeV, $m_t = 172.35235532883115$ GeV	declared D10/D11 split surface
Charged leptons	m_e, m_μ, m_τ witness values	target-anchored same-family witness
Selected-class quarks	exact running sextet on f_P	selected public quark frame
Neutrino absolute attachment	$0.017454720257976796, 0.019481987935919015, 0.05307522145074924$ eV	weighted-cycle absolute branch
Hadrons	source-only strong-binding descent plus empirical $e^+e^- \to$ hadrons closure surface	OPH hadron backend for source-only masses; measured hadron data for empirical closure rows

Exact non-hadron hit surface. If one asks only for exact hits, the paper-facing non-hadron bundle splits by lane as follows:

Concretely, the exact values shown on this surface are

$$\alpha^{-1}(0) = 137.035999177(21), \qquad P \simeq 1.6309682094,$$

$$(M_W, M_Z, m_H) = (80.377,\ 91.18797809193725,\ 125.1995304097179)\,\text{GeV},$$

$$(m_e, m_\mu, m_\tau) = (0.00051099895,\ 0.1056583755,\ 1.7769324651340912)\,\text{GeV},$$

$$(u, d, s, c, b, t) = (0.00216,\ 0.00470,\ 0.0935,\ 1.273,\ 4.183,\ 172.3523553288311)\,\text{GeV},$$

which matches the official PDG 2025 API running-quark target surface exactly, with the top coordinate taken from the PDG cross-section mass entry. The auxiliary direct-top row is a compare-only codomain with a corpus-limited no-go boundary. On the declared common-refinement transport-frame carrier, the restricted exact chain goes further and emits explicit exact forward Yukawas Y_u and Y_d. and

$$(m_1, m_2, m_3) = (0.017454720257976796,\ 0.019481987935919015,\ 0.05307522145074924)\,\text{eV},$$

with emitted weighted-cycle absolute splittings

$$\begin{aligned}\Delta m_{21}^2 &= 7.488059465106851 \times 10^{-5}\,\text{eV}^2,\\ \Delta m_{31}^2 &= 2.5123118727618473 \times 10^{-3}\,\text{eV}^2,\\ \Delta m_{32}^2 &= 2.4374312781107786 \times 10^{-3}\,\text{eV}^2.\end{aligned}$$

The table is lane-based. It states the derivation chain that emits each exact-hit surface and the exact caveat that prevents silent promotion to a stronger theorem claim.

How to read mismatches and exact hits. Exact numerical agreement is not by itself a proof of a blind prediction. The provenance ledger classifies the W/Z row as target-used frozen-reference reproduction, the charged-lepton triple as a target-anchored current-family witness, and the quark sextet as selected-class exact closure on the public frame class rather than a global classification of all quark frames. Where a row does not match, or where it matches only on a constrained surface, the explanations are as follows. The P-closure and $\alpha(0)$ displacement is localized to the Ward-projected Thomson residual. At the implemented fixed point the residual is

$$0.0411640123783505420500054142128067 38$$

inverse-alpha units. At the public endpoint pixel value

$$P \simeq 1.6309682094$$

the endpoint residual package requires

$$\Delta_{\text{source}}(P) = 0.04146586100522338905344871535731 4044\ldots.$$

That scalar belongs to the source-side hadronic spectral transport and scheme-remainder map. The OPH plus empirical hadron closure surface supplies the display row through the measured Thomson endpoint while the $e^+e^- \to$ hadrons payload class records the data-driven hadron path. The source-only transport row excludes an inserted comparison endpoint. The electroweak status table records the residual map and the RG/matching/threshold/scheme packet. The electroweak W/Z row is therefore compare-only. Charged leptons carry a corpus-limited no-go because the determinant trace-lift attachment from the D10 descendants of P to physical charged data is absent. The auxiliary direct-top PDG row differs from the theorem coordinate by 0.20673301656674425 GeV, or 0.28458848947515303 combined standard deviations, because Q007TP and Q007TP4 are distinct extraction codomains; the direct-top response kernel has a corpus-limited no-go boundary. Neutrino absolute masses are not directly measured, and PMNS-angle residuals remain visible comparison tension outside the theorem branch. Source-only hadron masses have no emitted prediction because the required Ward-projected hadronic spectral measure must come from a working OPH hadron backend. Empirical hadron closure rows carry a separate row class.

Local unification surface. The bosonic rows carry one cross-lane statement. The same pixel input P, fixed on the synthesis-paper outer/inner closure relation, fixes the D10/D11 bosonic trunk

$$P \longmapsto \alpha_U(P) \longmapsto (t_U(P), t_{\text{tr}}(P)) \longmapsto v(P) \longmapsto (M_W, M_Z)\,.$$

$$\sigma_{D11,\text{HT}} = \alpha_U(P)\cos(2\theta_{W0})/\sqrt{\pi} \longmapsto (m_H, m_t)\,.$$

Here $\alpha_U(P)$ is the branch value selected by the same forward D10 pixel-closure solve, so the bosonic trunk contains no inverse electroweak readback of the internal transmutation data. On the companion gravity side, the same pixel law packages

$$\bar{\ell}_{\text{SU}(2)}(t_{2,\text{run}}) \;+\; \bar{\ell}_{\text{SU}(3)}(t_{3,\text{run}}) = P/4, \qquad G = \frac{a_{\text{cell}}}{4\,\bar{\ell}(t)}.$$

The local unification surface separates one explicit familiar-unit readout package from three status surfaces: the W/Z compare-only validation sidecar, the exact Higgs/top theorem surface, and the gravity-side release with its strict classical-regime clause. On the gravity side, the stated local extension surface uses the lifted product presentation of the realized quotient branch and identifies

$$\bar{\ell}_{\text{shared}} = \bar{\ell}_{\text{SU}(2)}(t_{2,\text{run}}) \;+\; \bar{\ell}_{\text{SU}(3)}(t_{3,\text{run}}).$$

On that same surface the D10 pixel law fixes $\bar{\ell}_{\text{shared}} = P/4$, and the local SI readout is

$$G_{\text{SI}} = \frac{c^3 a_{\text{cell}}}{\hbar P}$$

relative to the declared microscopic datum a_{cell}. On that declared extension surface the same familiar-unit package reads

$$L_{\text{loc}} = \sqrt{a_{\text{cell}}}\,\widehat{L}(P), \qquad t_{\text{loc}} = \frac{\sqrt{a_{\text{cell}}}}{c}\,\widehat{T}(P),$$

$$E_{\rm loc} = \frac{\hbar c}{\sqrt{a_{\rm cell}}}\,\widehat{E}(P), \qquad \Theta_{\rm loc} = \frac{\hbar c}{k_B\sqrt{a_{\rm cell}}}\,\widehat{\Theta}(P),$$

with dimensionless $\widehat{L}, \widehat{T}, \widehat{E}, \widehat{\Theta}$. Thus, at fixed P, the local ruler is $\sqrt{a_{\rm cell}}$; seconds are that ruler divided by the structural Lorentz output c; and GeV and Kelvin are downstream familiar-unit displays of the inverse local ruler through $\hbar$ and k_B. On that declared extension surface the local release bundle is

$$c = 299792458\,\text{m/s}, \qquad G = 6.674299995910528 \times 10^{-11}\,\text{m}^3\,\text{kg}^{-1}\,\text{s}^{-2},$$

$$(M_W, M_Z, M_H) = (80.377,\ 91.18797809193725,\ 125.1995304097179)\,\text{GeV},$$

where M_W and M_Z are compare-only frozen-adapter values and M_H is the declared Higgs/top-surface value. The paper keeps c as the structural Lorentz output and keeps the G row as an exact emitted branch value relative to the stated $a_{\rm cell}$ datum, with no literal zero-difference identity against the rounded benchmark 6.6743×10^{-11}.

The detailed inventory below records the public theorem/continuation rows sector by sector. Whenever an exact sidecar or same-family witness is stronger than the public theorem row, the caveat is the one stated in the exact-hit table above.

3.1 Detailed particle inventory

Family	Particle	Claim tier	OPH value or strongest derived benchmark	Remaining step
Structural carriers	photon	structural	0	none
Structural carriers	gluons	structural	0 for the color gauge sector	none
Structural carriers	graviton	structural	0 for the dynamical-metric spin-2 carrier	none
Electroweak D10 branch	W boson	compare-only adapter	80.377 GeV on the frozen W/Z validation surface	compare-only validation beneath the declared D10/P/fine-structure arithmetic bridge; public validation records are source spectral measure payload, same-scheme remainder, and interval certificate
Electroweak D10 branch	Z boson	compare-only adapter	91.18797809193725 GeV on the frozen W/Z validation surface	same compare-only validation caveat as the W row
Higgs/top critical stage	Higgs boson	quantitative theorem	125.1995304097179 GeV on the declared D10/D11 running, matching, and threshold surface	exact source-only Higgs/top split theorem; compare-only inverse Higgs/top slice is on the same Jacobian surface
Quark family	top quark	selected-class theorem	172.35235532883115 GeV on the selected public quark frame class f_P	exact running-mass theorem row on f_P; the D11 lane carries the same exact top value on its declared split quantitative surface, while the repo-wide exact public top row is emitted by the selected-class quark theorem
Quark family	up quark	selected-class theorem	0.00216 GeV on the selected public quark frame class f_P	exact running-mass theorem row on f_P; supporting exact same-family and restricted transport-frame surfaces realize the same value; global classification of all public quark frame classes is closed as a corpus-limited no-go
Quark family	down quark	selected-class theorem	0.00470 GeV on the selected public quark frame class f_P	same selected-class theorem surface; target-free mass bridge on the emitted D12 ray closes separately
Quark family	strange quark	selected-class theorem	0.0935 GeV on the selected public quark frame class f_P	same selected-class theorem surface; the lower selector $\sigma_{\rm ref}$ is a negative sheet statement beneath that theorem
Quark family	charm quark	selected-class theorem	1.273 GeV on the selected public quark frame class f_P	same selected-class theorem surface; explicit exact forward Yukawas Y_u, Y_d are emitted on f_P
Quark family	bottom quark	selected-class theorem	4.183 GeV on the selected public quark frame class f_P	same selected-class theorem surface; the exact sextet matches the official PDG 2025 API running-quark target surface

Family	Particle	Claim tier	OPH value or strongest derived benchmark	Remaining step
Charged leptons	electron	continuation gap	n/a; exact same-carrier centered readback exists once a charged source pair is emitted, but the absolute scale is blocked by the closed common-shift no-go; compare-only target $g_e^\star = 0.0457789$, equivalently $\Delta_e^{\mathrm{abs},\star} = 3.00398633$	promote the latent candidate $\widehat{C}_e^{\mathrm{cand}}$ by closing branch-generator splitting, then emit the post-promotion lift whose descended scalar is $\mu_{\mathrm{phys}}(Y_e)$; within that lift the exact smaller forcing object is the physical identity-mode equalizer, after which $\widetilde{C}_e(Y_e) = \widehat{C}_e(Y_e) + \mu_{\mathrm{phys}}(Y_e)\,\mathbf{1}$, $A_{\mathrm{ch}}(Y_e) = \mu_{\mathrm{phys}}(Y_e)$, and the readouts g_e, Δ_e^{abs} follow
Charged leptons	muon	continuation gap	n/a; same exact centered-readback / common-shift-no-go frontier as the electron row	same $\widehat{C}_e^{\mathrm{cand}} \to$ promotion $\to$ post-promotion lift $\to \mu_{\mathrm{phys}}(Y_e) \to A_{\mathrm{ch}} \to g_e$ closure chain, with the physical identity-mode equalizer beneath the descended scalar
Charged leptons	tau lepton	continuation gap	n/a; same exact centered-readback / common-shift-no-go frontier as the electron row	same $\widehat{C}_e^{\mathrm{cand}} \to$ promotion $\to \mu_{\mathrm{phys}}(Y_e) \to A_{\mathrm{ch}} \to g_e$ closure chain, with the physical identity-mode equalizer beneath the descended scalar
Neutrinos	ν_e	weighted-cycle theorem branch	0.017454720257976796 eV; the weighted-cycle branch emits $\theta_{12} = 34.2259^\circ$, $\theta_{23} = 49.7228^\circ$, $\theta_{13} = 8.68636^\circ$, $\delta = 305.581^\circ$, $J = -0.02753$, $\Delta m^2_{21}/\Delta m^2_{32} = 0.03072111$, $\alpha_{21}^{(\mathrm{Maj})} = 153.618518^\circ$, $\alpha_{31}^{(\mathrm{Maj})} = 257.003241^\circ$, the bridge invariant $C_\nu = 0.9994295999075177$, the emitted proxy $P_\nu = 6.699825740519345$, the paper-facing amplitude $B_\nu = 6.696004159297337$, and $\lambda_\nu = 1.7237014208357415$	closed on the declared weighted-cycle theorem branch; the positive-segment adapter, bridge corridor, and bridge-coordinate sidecars are diagnostic-only and do not feed back into theorem state
Neutrinos	ν_μ	weighted-cycle theorem branch	0.019481987935919015 eV; same weighted-cycle theorem branch as the electron-neutrino row	same bridge-invariant / induced-amplitude package and same diagnostic-only retirement rule for the positive-segment adapter and bridge corridor
Neutrinos	ν_τ	weighted-cycle theorem branch	0.05307522145074924 eV; same weighted-cycle theorem branch as the electron-neutrino row	same bridge-invariant / induced-amplitude package and same diagnostic-only retirement rule for the positive-segment adapter and bridge corridor
Hadrons	proton	source backend absent; empirical closure policy emitted	no source-only emitted prediction	source-only prediction requires a working OPH hadron backend, such as GLORB/Echosahedron, plus a Ward-projected spectral-measure export and production systematics; empirical closure rows use a separate $e^+e^- \to$ hadrons payload class
Hadrons	neutron	source backend absent; empirical closure policy emitted	no source-only emitted prediction	same OPH-backend gate, with isospin-resolved hadron systematics in the backend row class
Hadrons	neutral pion proxy	source backend absent; empirical closure policy emitted	no source-only emitted prediction	same OPH-backend gate; local stable-channel surrogate output is not a paper prediction
Hadrons	$\rho(770)^0$ proxy	source backend absent; empirical closure policy emitted	no source-only emitted prediction	same OPH-backend gate, plus finite-volume resonance extraction in the backend row class

4 Detailed Closure Values, Screen Architecture, and Theorem Packages

This section records the detailed closure-value and theorem checklist used by the particle derivation: the five axioms, the screen-capacity closure, the local pixel closure, the regulated screen realization, and the imported OPH theorem packages. The order is the logical order used later in the paper: axioms and screen language first, closure values next, and theorem packages after that.

4.1 Canonical OPH basis

The particle-spectrum derivation uses the same five axioms as the compact reconstruction paper [4]. They are restated here for local readability because every later particle-family chapter depends on them either directly or through the structural OPH chain imported from that paper.

Axiom 4.1 (Screen Net)**.** *Physical data is organized on a horizon screen S^2 carrying a net of local algebras*

$$P \longmapsto \mathcal{A}(P)$$

for connected patches $P \subset S^2$, with isotony

$$P \subset Q \implies \mathcal{A}(P) \subset \mathcal{A}(Q).$$

Axiom 4.2 (Overlap Consistency)**.** *For overlapping patches $P_1 \cap P_2 \neq \varnothing$, the local states induced on the shared algebra agree:*

$$\omega_{P_1}|_{\mathcal{A}(P_1 \cap P_2)} = \omega_{P_2}|_{\mathcal{A}(P_1 \cap P_2)}.$$

Axiom 4.3 (Local MaxEnt and Refinement Stability)**.** *At the regulator scale ℓ_{UV}, the realized branch is selected by maximizing entropy subject to a finite family of gauge-invariant local constraints*

$$\mathcal{C}_{\ell_{\mathrm{UV}}} = \{O_a(x)\}_{a=1}^{N_{\mathrm{con}}},$$

with support radius $O(\ell_{\mathrm{UV}})$ and with label set independent of the number of regulator cells. Under refinement, the same finite constraint family persists, so the realized states belong to one common finite-dimensional MaxEnt family rather than unrelated maximizers.

Axiom 4.4 (Recoverable Generalized Entropy)**.** *A generalized entropy functional exists on caps,*

$$S_{\mathrm{gen}}(C) = S_{\mathrm{bulk}}(C) + \langle L_C \rangle,$$

where L_C is a positive edge-center entropy functional and the semiclassical branch identifies its leading coarse-grained contribution with $A(\partial C)/(4G)$. The functional obeys the recoverability and focusing structure required by the collar and null-modular arguments.

Axiom 4.5 (Minimal Admissible Realization)**.** *Among admissible realized low-energy sector packages $\mathfrak{S}$ consisting of the connected Lie gauge-sector image relevant in the EFT branch, its admissible light chiral matter content, and one Higgs doublet, the realized package is lexicographically minimal under*

$$C(\mathfrak{S}) = (\chi_{\mathrm{cpl}}, N_{\mathrm{nonab}}, N_c, N_g),$$

subject to loop coherence, anomaly freedom, refinement-stable light chiral matter, single-Higgs Yukawa completable structure with one connected abelian charge factor, intrinsic CP capability, and weak-sector UV completeness.

For particle physics, the first four axioms provide the observer-centric kinematic and entropic/modular background. The fifth axiom, MAR, is the selector that turns "some compact gauge group reconstructed from the transportable edge-sector category under the compact paper's explicit coherent refinement ladder, symmetry, and finite-fiber conditions" into the realized Standard Model branch. This paper therefore uses MAR constantly, but only after the compact-gauge reconstruction step has been separated from the later realized-branch selection.

4.2 Regulated screen architecture and patch language

The SM/GR derivation in Ref. [4] states the axioms in algebraic language. Ref. [3] supplies one concrete regulated realization of that language: a federation of finite observer patches with echosahedral multi-port interfaces, recurrent toroidal subchannels, exposed overlap packets, record algebras, and local repair instruments. A spherical screen remains a support-visible regulator chart, not a literal microscopic computer. This regulated architecture is important for the particle derivation for three reasons.

First, it makes the screen picture operational. A patch is a finite local algebra in a bounded carrier, and an overlap is a boundary-visible algebra with declared shared observables. Second, it makes the edge-sector language concrete. The particle derivation ultimately uses edge sectors, transport, refinement, fusion, and overlap data to build the gauge and flavor branches; Ref. [3] shows how those objects can be realized at fixed cutoff by finite gauge-aware patch carriers. Third, it clarifies the measurement interface. The observer-facing measurement package is a theorem-bearing fixed-cutoff statement about a central record algebra, Born probabilities for its event projectors, and Lüders conditioning on that same commuting algebra.

This regulated architecture should be read as a reference implementation, not as the unique OPH UV completion. Ref. [3] is explicit on that boundary: it selects a federated patch-carrier architecture because it is local, finite-dimensional, gauge-aware, naturally compatible with collars, overlap observables, and repair maps, and no longer asks the reader to identify the spherical regulator chart with the literal substrate. Hardware evidence is imported only when it appears in a public OPH evidence bundle with stable hashes and verifier receipts. This paper inherits that architecture as a concrete screen-side realization of the OPH language, while keeping microscopic uniqueness separate from the support-visible continuum theorem surfaces supplied by the compact paper.

4.3 Quantitative inputs and where they enter the particle derivation

The quantitative derivation developed here uses one Phase-II closure variable together with one external continuous input:

$$P \equiv a_{\text{cell}}/\ell_P^2, \tag{1}$$

$$N_{\text{scr}} \equiv \log \dim \mathcal{H}_{\text{tot}}. \tag{2}$$

The particle-spectrum draft needs to keep these closure values visible because several later sections depend on them in different ways.

Pixel area P. The pixel area is the declared local UV-area ratio. In this derivation it feeds the heat-kernel / edge-law input stage and, through that route, the forward D10 electroweak surface. It is the common upstream numerical variable for the electroweak, Higgs/top, flavor, quark, and hadron-facing branches. The particle paper therefore treats $P = P_\star$ as the inherited outer/inner fixed point

$$P_\star = \varphi + \frac{\sqrt{\pi}}{A_T(P_\star)}$$

from the synthesis paper [1], not as a quantity derived again inside the particle-spectrum argument itself.

Screen capacity N_{CRC}. The screen capacity is not part of the local recovered-core gravity/gauge derivation. It enters the separate screen-capacity branch through the cosmic record-closure target

$$N_{\mathrm{CRC}} = F(N_{\mathrm{CRC}}), \qquad \Lambda_{\mathrm{CRC}} = \frac{3\pi}{GN_{\mathrm{CRC}}}.$$

The count-density representation is

$$N_\star = \mathrm{MAR}\,\arg\max_N \left[\log|\Omega_N^{\mathrm{sc}}| - N\right].$$

In differentiable form, with $\ell(N) = \log|\Omega_N^{\mathrm{sc}}| - N$, this same branch is the unique stable fixed point of $T_\eta(N) = N + \eta\ell'(N)$ when the OPH normal-form count gives a strictly concave normalized density. Thus the capacity is the unique OPH-derived global readback balance point, not a separately fitted cosmological number. In the particle context, this matters mainly for a legacy capacity-level neutrino side estimate and for the cosmological-capacity discussion that frames why local null data do not determine the cosmological constant. The implemented branch still uses the static-patch normalization

$$N_{\mathrm{patch}} = \left(\frac{r_{\mathrm{dS}}}{\ell_P}\right)^2 \simeq 1.05 \times 10^{122}, \qquad N_{\mathrm{scr}} = \pi N_{\mathrm{patch}} \simeq 3.31 \times 10^{122}.$$

This paper should therefore keep the neutrino chapter direct: legacy capacity-level order-of-magnitude bookkeeping and the weighted-cycle neutrino derivation are not the same thing.

Why these are the only declared continuous closure values. One of the discipline constraints of the OPH derivation is that particle-sector freedom should not be hidden in a large uncontrolled parameter set. This paper exposes the two continuous closure values, $P_\star$ and the N_{CRC} readback fixed point, with $N_\star$ as its count representation, then asks the overlap, modular, gauge, and admissibility machinery to do the rest. Whether later continuation branches stay faithful to that discipline is one of the central questions this paper is meant to make explicit. For the particle-spectrum branch specifically, the nontrivial common quantitative burden sits on $P_\star$: N_{scr} enters the separate capacity/cosmology side, whereas the reverse-engineering claim for masses and couplings is that one shared pixel fixed point should drive the spectrum instead of a long sector-by-sector input list.

4.4 Technical theorem data carried into the particle-spectrum derivation

This paper uses the same theorem checklist as *Recovering Relativity and Standard Model Structure from Observer Overlap Consistency* [4], together with the regulator-language clarification integrated into the consensus and microphysics appendices. The ledger matters only insofar as it separates branch-internal statements from declared inputs.

The third axiom internalizes several pieces of infrastructure that often appear as separate assumptions: the local Gibbs form of the regulator-scale state, the quasi-local propagation bound, the endpoint-control estimate for bounded intervals, and the meaning of refinement stability itself. In regulator language, one works with finite local Hilbert spaces and a boundary-fixed overlap action on cut data, often summarized as the R0/R1 presentation. Beyond that regulator package, the compact paper supplies the support-visible BW scaling theorem for the Lorentz/null-modular/Einstein branch. Transportability, the fixed-cutoff bosonic sector category, the refinement/fiber ladder, and the realized MAR-admissible compact-gauge witness data are theorem-produced in the compact paper. On the central branch the matching gluing obstruction is the zero-obstruction transport criterion, not a second theorem-side input. Any local-Gibbs or collar-mixing language on that branch is fixed-cutoff recoverability/support control; it is not the source of the compact-gauge witness. For the BW/geometric side, the target is the support-visible extracted prime geometric subnet. The compact paper proves the needed scaling theorem by combining regularized support-visible modular transport, weak-$*$/GNS extraction of the cap pair, support-readable modular covariance, ordered cut-pair rigidity, and KMS/BW normalization. The unregularized full-algebra common-floor route is deliberately not claimed, because off-support directions may collapse without affecting observer-facing matrix elements.

For the particle-spectrum derivation, this has a practical consequence. Structural statements such as compact-gauge reconstruction, the realized Standard Model quotient, the exact hypercharge lattice, the color count $N_c = 3$, the generation count $N_g = 3$, and the symmetry-protected massless carrier zeros are not floating independently of the declared input and theorem checklist. They live on a controlled chain whose scope conditions need to be visible in this paper, especially once later chapters turn from structural branches to quantitative closure, continuation, and nonperturbative sectors.

4.5 Core theorem packages used later

This paper will repeatedly use earlier OPH results without re-proving them in full. The theorem packages imported into the present draft are the following.

1. **Patch-net and overlap language.** Ref. [2] organizes overlap repair, fixed-point language, cycle holonomy, and gauge-as-gluing in a form that is especially useful for flavor transport and observer-centric interpretation.

2. **Screen-side regulated architecture.** Ref. [3] gives the federated echosahedral patch-carrier model, the regulated patch-net embedding theorem, and the fixed-cutoff edge heat-kernel / Casimir theorem.

3. **Measurement interface.** Ref. [3], together with the integrated measurement appendices and the synthesis paper *Observers Are All You Need* [1], supplies the fixed-cutoff central-record, Born-rule, and Lüders-conditioning package together with the fixed-cutoff Bell/CHSH theorem stack used when later chapters discuss how measurements pick out definite observer-accessible outcomes and how two-wing comparison laws stay on the quantum side of the classical Bell bound.

4. **Compact gauge reconstruction and Standard Model structure.** *Recovering Relativity and Standard Model Structure from Observer Overlap Consistency* [4] and the dedicated gauge-group fragment together provide the path from refinement-stable edge sectors to a compact group, then from MAR to

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_g = 3, \qquad N_c = 3.$$

5. **Continuation-level topic notes.** The integrated lane appendices carry useful material on supersymmetry, heuristic baryogenesis continuations, proton stability, generation structure, and Yukawa hierarchy, but those sections have to be read with their stated status boundaries intact. They are sources for later discussion and continuation chapters, not automatic upgrades to recovered-core theorems.

This reading rule governs the whole paper. Structural gauge and carrier results can be used as structural results. The electroweak branch is a forward-emitted Phase-II quantitative branch. The Higgs/top critical stage is a source-only split quantitative theorem on its declared running, matching, and threshold surface. The quark, charged-lepton, neutrino, and hadron chapters retain their stated continuation or simulation status unless the underlying derivation chains close more strongly.

5 Gauge Architecture, Standard Model Structure, and Structural Carriers

The particle-spectrum derivation does not start from a list of particles. It starts from the compact-gauge branch. OPH reconstructs a compact gauge group from the zero-obstruction edge-sector category produced by the compact paper's transportability, fixed-cutoff category, and refinement/fiber theorems. The overlap/transport obstruction calculus is the classification stage, not the selection stage. The compact-gauge witness theorem supplies nonempty realized MAR-admissible branch data, and MAR then selects the realized low-energy package. That logical order matters for this paper. The gauge branch fixes the structural carrier content before any mass readout, and it also explains why some rows are exact structural zeros rather than fitted or accidentally small numbers.

On the realized branch, the gauge/content result used throughout this paper is

$$\frac{\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)}{\mathbb{Z}_6}, \qquad N_g = 3, \qquad N_c = 3.$$

Recovering Relativity and Standard Model Structure from Observer Overlap Consistency [4], the longer main derivation surface, and the dedicated gauge-group proof fragment all agree on the logical split: compact-gauge reconstruction first, realized Standard Model selection second. The first stage gives a compact internal symmetry group under the stated zero-obstruction bosonic sector conditions. The second stage uses MAR and anomaly cancellation to fix the realized quotient and hypercharge lattice; the minimal coupled carrier fixes $N_c = 3$; intrinsic CP capability together with weak-sector UV completeness bounds N_g; MAR then fixes $N_g = 3$; and Witten parity remains a consistency check on the realized triplet-doublet package.

For the particle zoo this structural branch fixes three exact carrier statements. First, the photon is massless on the realized electromagnetic branch. Second, the gluons are massless on the color branch, even though free gluons are never seen as asymptotic particles because color is confined. Third, the graviton is massless on the dynamical-metric branch of the gravity chain. These are symmetry-protected zeros of the realized branch, not continuation-level masses.

These structural zeros should not be confused with the later quantitative particle rows. The paper should therefore treat the photon, gluons, and graviton as consequences of the gauge and gravity branches themselves, not as outputs of the later quantitative or flavor machinery. The later boson sections will add the W boson, the Z boson, and the Higgs boson, but those are not structural zeros; they sit on electroweak and Higgs/top quantitative stages.

6 Electroweak D10 Branch and the Higgs/Top Critical Stage

The reported bosonic sector distinguishes sharply between emitted numerical rows and scoped non-emitted rows. The active bosonic stages are the electroweak D10 prediction stage and the Higgs/top critical stage, while charged-lepton rows carry their no-go boundary, the neutrino lane closes on a separate weighted-cycle theorem branch, and hadron rows are backend-gated. The bosonic numerical sector is therefore concentrated in two places: the D10 electroweak branch for the W boson and Z boson, and the Higgs/top critical stage for the Higgs boson together with the exact D11 split-theorem top value.

6.1 Electroweak prediction stage on the D10 quantitative-closure branch

The electroweak derivation consists of a single-P running family followed by a reduced two-scalar carrier, a selector on that carrier, an exact carrier mass chart, and then a target-free source-only transport theorem beyond the selected carrier. In practical terms, the construction starts from the declared pixel input P, builds the running electroweak family, reduces it to the selected carrier, reads the W boson/Z boson pair from the selected point on that carrier, and then emits the mass pair together with the Ward-projected electromagnetic transport family from the source basis alone. On the D10 branch, a live forward prediction must respect exactly that ordering: first certify the shared pixel input P on the declared D10 running/matching/threshold/scheme surface, thereby fixing the source basis

$$(\alpha_{2,m_Z}, \alpha_{Y,m_Z}, \eta_{\text{source}}, v),$$

with $\alpha_U(P)$, equivalently $t_U(P)$ and $t_{\text{tr}}(P)$, fixed by the same forward pixel-closure solve, and only then emit the electroweak transport family forward from that source basis. The completion pipeline classifies that full chain outside the promoted theorem tier. Its W/Z rows are the frozen validation adapter values, and their exact numerical agreement is recorded as comparison status rather than as a closed D10 Phase-II prediction theorem.

The same quantitative lane has one exact golden-ratio benchmark. If one writes

$$x(C) := \frac{S_{\text{gen}}(C)}{S_{\text{bulk}}(C)} = 1 + \frac{\langle L_C \rangle}{S_{\text{bulk}}(C)}$$

for the total/bulk/edge hierarchy and imposes exact self-similar balance

$$\frac{S_{\text{gen}}(C)}{S_{\text{bulk}}(C)} = \frac{S_{\text{bulk}}(C)}{\langle L_C \rangle},$$

then x obeys $x^2 - x - 1 = 0$, so the unique positive equilibrium point is

$$x = \varphi := \frac{1 + \sqrt{5}}{2}.$$

Equivalently the order parameter $A_\varphi(x) = x - 1 - \frac{1}{x}$ vanishes there. The synthesis paper fixes the exact numerical value of P through the outer/inner closure relation. The point of the equilibrium theorem here is that the proximity to φ has a structural reason. Exact balance is too symmetric to support durable records, structure, and dynamics, so the closure solve fixes a small equilibrium-breaking detuning away from that balance point.

The compare-only W/Z adapter rows in the final bundle are

$$m_W = 80.377 \text{ GeV}, \qquad m_Z = 91.18797809193725 \text{ GeV}.$$

These numbers sit on the frozen validation surface, not on a target-free prediction theorem. The exact selected-carrier chart is explicit on disk and emits the local pair

$$m_W^{\text{carrier}} = 80.38629169244275 \text{ GeV}, \qquad m_Z^{\text{carrier}} = 91.18290444674243 \text{ GeV},$$

so the mathematics distinguishes three electroweak objects: the exact selected-carrier chart, the freeze-once coherent repair surface, and the target-free source-only theorem surface that would supply public prediction rows once the P, Thomson, and RG/matching/threshold/scheme gates are closed.

Theorem 6.1 (Freeze-once coherent electroweak repair law)**.** *Let the emitted running/core electroweak basis be*

$$Q_{\rm run}(P) = \big(\alpha_{Y,m_Z}(P), \alpha_{2,m_Z}(P), v(P), \eta_{\rm source}(P)\big),$$

and write

$$\tau_Y(\tau_2) = -\frac{\tau_2 + 2\eta_{\rm source}}{1+4\tau_2^2}, \qquad n_{\rm fiber}(\tau_2) = 1 + \frac{\alpha_{Y,m_Z}\tau_Y(\tau_2) + \alpha_{2,m_Z}\tau_2}{\alpha_{Y,m_Z} + \alpha_{2,m_Z}}.$$

Freeze one authoritative target pair

$$\mathcal{T}^\dagger = (M_W^\dagger, M_Z^\dagger), \qquad 0 < M_W^\dagger < M_Z^\dagger,$$

and define the charged anchor

$$\tau_{2,W}^\dagger = \frac{M_W^{\dagger\,2}}{\pi v^2 \alpha_{2,m_Z}} - 1, \qquad \delta\alpha_2^\dagger = \frac{M_W^{\dagger\,2}}{\pi v^2} - \alpha_{2,m_Z}.$$

Then the fiber-parallel hypercharge leg is

$$\tau_Y^\dagger = \tau_Y(\tau_{2,W}^\dagger), \qquad n_{\rm fiber}^\dagger = n_{\rm fiber}(\tau_{2,W}^\dagger),$$

$$M_{Z,\rm fiber}^\dagger = v\sqrt{\pi(\alpha_{Y,m_Z} + \alpha_{2,m_Z})\, n_{\rm fiber}^\dagger},$$

and the orthogonal neutral-shear scalar is

$$\delta M_Z^\perp = M_Z^\dagger - M_{Z,\rm fiber}^\dagger,$$

equivalently

$$\delta n^\dagger = \frac{(M_Z^\dagger + M_{Z,\rm fiber}^\dagger)\,\delta M_Z^\perp}{\pi v^2(\alpha_{Y,m_Z} + \alpha_{2,m_Z})}, \qquad \delta\alpha_Y^\perp = \frac{(M_Z^\dagger + M_{Z,\rm fiber}^\dagger)\,\delta M_Z^\perp}{\pi v^2}.$$

The fiber-parallel hypercharge motion is

$$\delta\alpha_Y^\parallel = \alpha_{Y,m_Z}\,\frac{8\eta_{\rm source}(\tau_{2,W}^\dagger)^2 - \tau_{2,W}^\dagger}{1 + 4(\tau_{2,W}^\dagger)^2}.$$

Therefore the unique coherent repair package is

$$\Sigma_{EW}^\dagger = (\delta\alpha_2^\dagger, \delta\alpha_Y^\parallel, \delta\alpha_Y^\perp),$$

equivalently $(\tau_{2,W}^\dagger, \delta n^\dagger)$. *If the selected carrier anchor is the* $\tau_2 = 0$ *fiber point, so that*

$$\alpha_{2,*} = \alpha_{2,m_Z}, \qquad \alpha_{Y,*} = \alpha_{Y,m_Z}(1 - 2\eta_{\rm source}),$$

then the coherent validation couplings are

$$\alpha_{2,\dagger} = \alpha_{2,*} + \delta\alpha_2^\dagger, \qquad \alpha_{Y,\dagger} = \alpha_{Y,*} + \delta\alpha_Y^\parallel + \delta\alpha_Y^\perp,$$

and they emit one coherent quintet

$$M_W^\dagger = v\sqrt{\pi\alpha_{2,\dagger}}, \qquad M_Z^\dagger = v\sqrt{\pi(\alpha_{Y,\dagger}+\alpha_{2,\dagger})},$$

$$\alpha_{\mathrm{em},\dagger}^{-1} = \frac{\alpha_{Y,\dagger}+\alpha_{2,\dagger}}{\alpha_{Y,\dagger}\alpha_{2,\dagger}}, \qquad \sin^2\theta_{W,\dagger} = \frac{\alpha_{Y,\dagger}}{\alpha_{Y,\dagger}+\alpha_{2,\dagger}}.$$

The frozen-target electroweak validation law is one unique coherent value-emission law on top of the closed fiber map.

For the local PDG-2025 comparison surface,

$$\begin{aligned} \alpha_{2,m_Z} &= 0.03377843630219015, & v &= 246.76711732749683\ \mathrm{GeV},\\ \alpha_{Y,m_Z} &= 0.010131601067241624, & M_W^\dagger &= 80.377\ \mathrm{GeV},\\ \eta_{\mathrm{source}} &= 0.022147000871961295, & M_Z^\dagger &= 91.18797809193725\ \mathrm{GeV}, \end{aligned}$$

and the law emits

$$\begin{aligned} \tau_{2,W}^\dagger &= -2.3116268316325517\times 10^{-4},\\ \delta\alpha_2^\dagger &= -7.808313968675484\times 10^{-6},\\ M_{Z,\mathrm{fiber}}^\dagger &= 91.17717018044121\ \mathrm{GeV},\\ \delta M_Z^\perp &= 10.807911496044653\ \mathrm{MeV},\\ \delta\alpha_Y^\parallel &= 2.3421435088707883\times 10^{-6},\\ \delta\alpha_Y^\perp &= 1.0302892666790446\times 10^{-5},\\ \delta n^\dagger &= 2.3463639031113336\times 10^{-4}. \end{aligned}$$

So the coherent validation couplings are

$$\alpha_{2,\dagger} = 0.03377062798822147, \qquad \alpha_{Y,\dagger} = 0.00969547694807616,$$

and the same coherent branch emits

$$M_W^\dagger = 80.377\ \mathrm{GeV}, \qquad M_Z^\dagger = 91.18797809193725\ \mathrm{GeV},$$

$$\alpha_{\mathrm{em},\dagger}^{-1} = 132.752408550967, \qquad \sin^2\theta_{W,\dagger} = 0.223058333896849.$$

These two carrier-side scalars are bookkeeping data on the frozen validation surface. They are not the public electromagnetic readout on the Ward-projected D10 lane. That freeze-once law provides an exact coherent validation surface. On that frozen authoritative validation surface the canonical W/Z pair is hit exactly, but only as compare-only validation. The intended public electroweak prediction surface is the target-free source-only transport law. Its public validation records are the source spectral measure payload, same-scheme remainder, and interval certificate. The displayed W/Z numbers below remain frozen-adapter comparison values.

Theorem 6.2 (Target-free source-only electroweak mass transport law)**.** *Let the D10 source-only emitted basis be*

$$Q_{\mathrm{src}}(P) = \big(\alpha_{2,m_Z}(P), \alpha_{Y,m_Z}(P), \eta_{\mathrm{source}}(P), v(P)\big),$$

and define

$$\rho_{\mathrm{EW}} := \frac{\alpha_{2,m_Z}-\alpha_{Y,m_Z}}{\alpha_{2,m_Z}+\alpha_{Y,m_Z}}, \qquad \alpha_u^{\mathrm{seed}} := \frac{\eta_{\mathrm{source}}}{\rho_{\mathrm{EW}}}, \qquad \lambda_{\mathrm{EW}} := \frac{\eta_{\mathrm{source}}\alpha_u^{\mathrm{seed}}}{4} = \frac{\eta_{\mathrm{source}}^2}{4\rho_{\mathrm{EW}}}.$$

Then the beyond-selected-carrier transport map is uniquely

$$\tau_2^{\text{exact}} = -\lambda_{\text{EW}}\left(1 + \frac{2}{3}\eta_{\text{source}} + \left(1 - \frac{\rho_{\text{EW}}}{6}\right)\eta_{\text{source}}^2\right),$$

$$\delta n^{\text{exact}} = \lambda_{\text{EW}}\left(1 + \frac{4}{3}\eta_{\text{source}} + \left(2 - \frac{\rho_{\text{EW}}}{6}\right)\eta_{\text{source}}^2\right),$$

with fiber hypercharge transport

$$\tau_Y^{\text{fiber}} = -\frac{\tau_2^{\text{exact}} + 2\eta_{\text{source}}}{1 + 4(\tau_2^{\text{exact}})^2}.$$

The coherent transport shifts are

$$\delta\alpha_2 = \alpha_{2,m_Z}\tau_2^{\text{exact}}, \qquad \delta\alpha_Y^{\parallel} = \alpha_{Y,m_Z}\frac{8\eta_{\text{source}}(\tau_2^{\text{exact}})^2 - \tau_2^{\text{exact}}}{1 + 4(\tau_2^{\text{exact}})^2},$$

$$\delta\alpha_Y^{\perp} = (\alpha_{2,m_Z} + \alpha_{Y,m_Z})\,\delta n^{\text{exact}}.$$

With

$$\alpha_{2,*} = \alpha_{2,m_Z}, \qquad \alpha_{Y,*} = \alpha_{Y,m_Z}(1 - 2\eta_{\text{source}}),$$

the coherent transport couplings are

$$\alpha_2' = \alpha_{2,*} + \delta\alpha_2, \qquad \alpha_Y' = \alpha_{Y,*} + \delta\alpha_Y^{\parallel} + \delta\alpha_Y^{\perp},$$

and they emit one coherent family

$$W = v\sqrt{\pi\alpha_2'}, \qquad Z = v\sqrt{\pi(\alpha_2' + \alpha_Y')}, \qquad v = v(P).$$

$$a_0(P) := \frac{\alpha_{2,m_Z}(P) + \alpha_{Y,m_Z}(P)}{\alpha_{2,m_Z}(P)\alpha_{Y,m_Z}(P)}, \qquad s_0(P) := \frac{\alpha_{Y,m_Z}(P)}{\alpha_{2,m_Z}(P) + \alpha_{Y,m_Z}(P)}.$$

So the target-free source-only transport law fixes the public mass pair together with the source-locked running-family anchor (a_0, s_0, v) *from the D10 basis alone. The physical electromagnetic row is read from the Ward-projected unbroken* $\mathrm{U}(1)_Q$ *channel in Theorem 6.3.*

Using the same frozen electroweak validation basis

$$\begin{aligned}
\alpha_{2,m_Z} &= 0.03377843630219015,\\
\alpha_{Y,m_Z} &= 0.010131601067241624,\\
\eta_{\text{source}} &= 0.022147000871961295,\\
v &= 246.76711732749683\ \text{GeV},
\end{aligned}$$

the target-free transport law would emit

$$W = 80.37700001539531\ \text{GeV}, \qquad Z = 91.18797807794321\ \text{GeV},$$

$$a_0 = 128.30576920234813, \qquad s_0 = 0.23073542347506173.$$

Theorem 6.3 (Ward-projected $\mathrm{U}(1)_Q$ transport law and Thomson-limit electromagnetic readout)**.** *Let the D10 source-only running family be locked by the forward pixel solve on the declared running/matching/threshold/scheme surface, with source basis*

$$Q_{\mathrm{src}}(P) = \big(\alpha_{2,m_Z}(P), \alpha_{Y,m_Z}(P), \eta_{\mathrm{source}}(P), v(P)\big).$$

Assume the realized D9 branch, so that

$$Q = T_3 + Y$$

on the realized low-energy branch and color-singlet physical states carry integer Q-charge. Assume further that the post-selector electroweak transport object is the populated quotient-local kernel

$$K^{\mathrm{EW}}_{D10}(q^2;P) = \{\Pi_{AA}(q^2;P), \Pi_{AZ}(q^2;P), \Pi_{ZZ}(q^2;P), \Pi_{WW}(q^2;P)\}$$

on the declared physical observable algebra, together with a Ward projector

$$\mathcal{W}_Q : K^{\mathrm{EW}}_{D10}(q^2;P) \longrightarrow \Pi_Q(q^2;P)$$

onto the unbroken electromagnetic channel satisfying at $q^2 = 0$

$$\mathcal{W}_Q[\Pi_{AZ}(0;P)] = 0, \qquad \mathcal{W}_Q[\Pi_{AA}(0;P)] \neq 0.$$

Assume the abelian projected edge-sector probabilities obey the $\mathrm{U}(1)$ *heat-kernel law*

$$p_n(q^2;P) \propto e^{-t_Q(q^2;P)\lambda_n}, \qquad \lambda_n = n^2,$$

equivalently

$$g_Q^2(q^2;P) = \frac{t_Q(q^2;P)}{2\pi},$$

and that the scalar readout package satisfies the exact provenance lock

$$\begin{aligned} \textit{family\_source\_id} &= \texttt{d10\_running\_tree}, \\ \textit{scheme\_id} &= \texttt{freeze\_once}, \\ \textit{origin\_kernel\_id} &= \texttt{EWTransportKernel\_D10}. \end{aligned}$$

Then the unique electromagnetic coupling readout on the Ward-projected D10 lane is

$$\alpha_{\mathrm{em}}^{-1}(q^2;P) = \frac{8\pi^2}{t_Q(q^2;P)}.$$

If

$$a_0(P) := \alpha_{\mathrm{em}}^{-1}(m_Z^2;P),$$

then the unique zero-normalized affine scalar on the same source-locked family is

$$\delta_\alpha(q^2;m_Z^2;P) := \frac{t_Q(m_Z^2;P)}{t_Q(q^2;P)} - 1,$$

so that

$$\alpha_{\mathrm{em}}^{-1}(q^2;P) = a_0(P)\big(1 + \delta_\alpha(q^2;m_Z^2;P)\big).$$

Hence the Thomson-limit electromagnetic readout on that same Ward-projected electromagnetic kernel:

$$\alpha_{\mathrm{Th}}^{-1}(P) := \lim_{q^2\to 0} \alpha_{\mathrm{em}}^{-1}(q^2;P) = a_0(P)\,\frac{t_Q(m_Z^2;P)}{t_Q(0;P)}.$$

Consequently, the low-energy electromagnetic row on this lane is read as the Thomson endpoint of the D10 electromagnetic transport family rather than as a mass-chart byproduct.

Proof. The realized D9 branch fixes the electric charge operator by $Q = T_3 + Y$, and the quotient-protected charge theorem fixes the physical color-singlet Q-lattice to be integral. The abelian edge-sector theorem supplies the heat-kernel extraction rule

$$p_n \propto e^{-t_Q \lambda_n}, \qquad g_Q^2 = \frac{t_Q}{2\pi},$$

so the Ward-projected Q-channel determines a unique transport time $t_Q(q^2; P)$. Because the Ward projector kills A - Z mixing at $q^2 = 0$, the selected channel is the physical unbroken electromagnetic channel. Then

$$\alpha_{\text{em}} = \frac{g_Q^2}{4\pi} = \frac{t_Q}{8\pi^2},$$

hence

$$\alpha_{\text{em}}^{-1} = \frac{8\pi^2}{t_Q}.$$

Evaluating at $q^2 = m_Z^2$ defines the source-locked anchor $a_0(P)$, and dividing by the same expression at general q^2 yields

$$\delta_\alpha = \frac{t_Q(m_Z^2)}{t_Q(q^2)} - 1.$$

The provenance-equality clause forces one running-family source, one frozen scheme, and one origin kernel for all scalar readouts, so no mixed-family or Z-only surrogate is promotable. The Thomson formula is the $q^2 \to 0$ limit of the same identity. □

Corollary 6.4 (Source-only release criterion on the Ward-projected D10 lane)**.** *The Ward-projected D10 lane supports a source-only fine-structure release row only when the populated Ward-projected transport kernel satisfies*

$$\lim_{q^2 \to 0} \alpha_{\text{em}}^{-1}(q^2; P) = \alpha_{\text{Th}}^{-1}(P),$$

on the same source-locked family and without any separate endpoint value feeding that lane.

The endpoint table records the source-locked anchor from the runtime candidate,

$$a_0(P) = 128.30826804516521389255200599018177893545 0\ldots,$$

and the OPH plus empirical hadron closure display value

$$\alpha_{\text{Th}}^{-1} = 137.035999177.$$

The repository's `code/P_derivation` witness path records the source-side audit trunk without a built-in reference inverse-α value. The public endpoint value is excluded as a source-only closure-solve input. The source-side audit trunk emits

$$\alpha_{\text{cand}}^{-1} = 136.9948351646216494579499945857871932620 29,$$

so the difference from the endpoint readout is recorded as an endpoint-and-matching audit packet.

Remark 6.5. *The empirical closure reference is the Thomson-limit inverse fine-structure constant. The source anchor* $a_0(P) = 128.308268045165213892552005990181778935450\ldots$ *is the electroweak-scale running-family value at* m_Z^2*, while the display row is the low-energy endpoint. The 2022 CODATA/NIST reference value is* $\alpha^{-1} = 137.035999177(21)$ *[5].*

The freeze-once coherent repair surface provides compare-only validation, with machine-scale deltas

$$\Delta M_W \approx 1.54\times10^{-8}\ \text{GeV},\qquad \Delta M_Z \approx -1.40\times10^{-8}\ \text{GeV}.$$

The companion carrier-side pair

$$\alpha^{-1}_{\text{em},\dagger} = 132.752408550967,\qquad \sin^2\theta_{W,\dagger} = 0.223058333896849$$

is mass-chart bookkeeping on that frozen validation surface. It is not the public electromagnetic readout on the Ward-projected $\mathrm{U}(1)_Q$ lane.

The electroweak quantitative-closure branch also contains several subordinate theorem objects: an underdetermination theorem for the quadratic repair family, the smallest theorem route through $\text{ColorBalancedQuadraticRepairDescent}_{EW}$, the source-only mass emitter beneath the public W/Z rows, and the Ward-projected electromagnetic transport theorem above the source-locked running-family anchor. The scalar-package consistency clause is the shared provenance lock

$$(\texttt{d10\_running\_tree}, \texttt{freeze\_once}, \texttt{EWTransportKernel\_D10}).$$

A further forward transmutation certificate sits directly beneath the promoted readout and records that the source-only basis reconstructs the same $\alpha_U(P)$, $t_U(P)$, and $t_{\text{tr}}(P)$ as the pixel-closure solve without reading them back from measured couplings. These objects locate the public repair law inside the quantitative-closure branch without changing the emitted public surface.

6.2 Higgs/top critical stage

The Higgs/top critical stage is a quantitative-theorem stage built on top of the electroweak gauge core. It is described by a source-only split law together with a Jacobian readout map on the declared D10/D11 running, matching, and threshold surface. The promoted D10 tuple

$$(\eta_{\text{source}}, \beta_{EW}, \lambda_{EW}, \tau^{\text{exact}}_{2,\text{tree}}, \delta n^{\text{exact}}_{\text{tree}})$$

emits the shared scalar

$$\rho_{HT} = \log\!\left(1+\tau^{\text{exact}}_{2,\text{tree}}\right)$$

and the source-only residual selectors

$$R_T = -\tau^{\text{exact}}_{2,\text{tree}}\eta^2_{\text{source}} + \left(1+\frac{\beta_{EW}}{28}\right)\eta^6_{\text{source}} + \frac{\eta^8_{\text{source}}}{14} + \frac{\eta^9_{\text{source}}}{27},$$

$$R_H = \eta^5_{\text{source}} - \frac{3}{25}\eta^6_{\text{source}} + \frac{\lambda_{EW}\eta^6_{\text{source}}}{18} + \frac{\eta^8_{\text{source}}}{2\beta_{EW}}.$$

The split coordinates are

$$\pi_y = \frac{\eta_{\text{source}} + \left(\frac{3}{2}+\frac{\beta_{EW}}{4}\right)\rho_{HT} + R_T}{\sqrt{\pi}},\qquad \pi_\lambda = \frac{\eta_{\text{source}} - \left(\frac{4}{3}-\frac{\beta_{EW}}{54}\right)\rho_{HT} + R_H}{\sqrt{\pi}},$$

so the declared D11 Jacobian reads out

$$\delta y_t(\mu_t) = \pi_y\, y_t^{\text{core}}(\mu_t),\qquad \delta\lambda(\mu_t) = -\frac{16}{9}\pi_\lambda\,\lambda^{\text{core}}(\mu_t).$$

This gives

$$m_H = 125.1995304097179\ \text{GeV},\qquad m_t^{D11} = 172.3523553288312\ \text{GeV}.$$

At the precision quoted by the PDG Higgs listing, the Higgs row lands on the 2025 Higgs average. The same surface emits a companion top coordinate. The exact public running-top row is carried by the selected-class quark theorem and uses the PDG 2025 cross-section entry. The bridge to the auxiliary direct-top PDG row is closed as a corpus-limited codomain no-go; the auxiliary row is compare-only. The one-scalar companion seed

$$\sigma_{D11,\mathrm{HT}} = \frac{\alpha_U \cos(2\theta_{W0})}{\sqrt{\pi}}$$

is on disk only as the fixed-ray companion branch with $\pi_y = \pi_\lambda = \sigma_{D11,\mathrm{HT}}$. Separately, the same D11 Jacobian admits a compare-only inverse slice that hits the canonical Higgs/top reference pair exactly. That exact sidecar is useful for bookkeeping and does not define the forward branch used for the public rows. The Higgs boson belongs naturally in the boson discussion, while the top quark will be revisited in the quark-family chapter as the third-generation up-type quark. The reported top-quark row here is the D11 split-theorem value on this surface, while the exact public top row on the repo-wide non-hadron mass surface is also carried by the selected-class quark theorem.

The distinction is between the closed source-only split theorem and the broader paper surface. The split theorem is closed on its emitted D10/D11 branch, so the Higgs/top rows are substantive emitted outputs. The exact inverse slice is a compare-only validation surface, and the theorem is confined to the declared D10/D11 quantitative surface rather than promoted into the recovered core or into the selected-class exact quark theorem.

The electroweak chapter therefore presents three mathematically distinct surfaces: the target-free source-only theorem surface that would supply promoted prediction rows after the declared RG/matching/threshold/scheme surface is internally certified, the selected-carrier chart beneath it, and the freeze-once coherent repair surface used for the compare-only W/Z values. That is the technically direct way to present the bosonic sector while keeping validation bookkeeping and matching/scheme gates separate from the declared P-closure and fine-structure endpoint.

7 Flavor Transport, Generation Structure, and the Yukawa Dictionary

The gauge branch explains why the realized low-energy sector has three generations, but it does not by itself produce the full flavor dictionary. This paper therefore needs a separate flavor chapter. The imported realized-branch input here is the D9 count package $N_c = 3$ and $N_g = 3$. This chapter is the bridge between that structural result and the matter-family chapters for quarks, charged leptons, and neutrinos.

The flavor derivation is deliberately constructive. It does not claim that the full OPH flavor observable is theorem-level. Instead, it builds the object chain that a closed theorem would have to pass through. The derivation chain has the following main architecture: start from a refinement-indexed family transport kernel, derive a centered generation-bundle branch generator, lift that to same-label transport data, derive the induced overlap-edge transport cocycle, reduce the cocycle to a persistent flavor observable, and then push that observable into common sector-response objects for the downstream quark, charged-lepton, and neutrino derivations.

Mathematically, this derivation is where "why this family exists" becomes a concrete technical question. The realized Standard Model branch gives three generations structurally; the flavor derivation is where one tries to turn that three-generation existence statement into transport, splitting, suppression, phase, and excitation data. The relevant objects are the intermediate transport and spectral structures, not the final fermion masses themselves, that later mass readouts consume.

The active chain, in slightly compressed form, is:

1. normalize a refinement-indexed family transport kernel;
2. derive a centered generation-bundle branch generator on the realized three-generation charged bundle;
3. lift this to same-label edge-line transport and then to an overlap-edge cocycle with explicit defect and gap bookkeeping;
4. reduce those data to projectors, spectral gaps, pair suppressions, and cycle phases;
5. push the resulting family object into sector-response objects and then into suppression/phase tensors for the downstream matter derivations.

Three lane-specific claim boundaries are attached to this region of the derivation. The shared excitation dictionary is the common proof-facing base. Above that base, the charged lane carries exact centered readback, a closed common-shift no-go, the declared same-label q_e readback, a source-side determinant character for any fixed formal exponent vector

$$S_M = \sum_e M_e^{\mathrm{ch}} \log q_e,$$

a determinant-line lift on theorem-grade physical charged data, and an algebraic mass readout from theorem-grade $A_{\mathrm{ch}}(P)$. The theorem lane does not emit a theorem-grade sector-isolated charged determinant exponent vector, and it does not identify a source-side determinant character with the physical charged determinant line. The charged-lepton theorem gap is the determinant trace-lift attachment $3\mu(r) = S_M(r)$ on the charged determinant channel. The charged determinant channel has a corpus-limited no-go boundary: the uncentered trace lift is not emitted. The neutrino lane closes on the weighted-cycle bridge-rigidity surface above the emitted proxy P_ν, with induced paper-facing amplitude B_ν and absolute attachment above the scale-free branch, and the quark lane carries the internalized mass bridge together with the selected-class public theorem on the public physical quark frame class f_P. That theorem makes the exact physical sigma datum target-free public on f_P, emits (g_u, g_d) algebraically through the affine mean law, and emits the exact running sextet together with explicit exact forward Yukawas Y_u and Y_d. A separate compare-only microphysical bridge records edge-statistics transport on a diagnostic surface. Together these boundaries mark where the flavor derivation leaves the common transport backbone and passes to lane-specific closure contracts.

8 Quark Family Derivation

The quark lane carries a theorem-grade selected-class closure on the public physical quark frame class f_P chosen by P. On that selected class the exact physical sigma datum is target-free public, the affine mean law emits (g_u, g_d) algebraically, the ordered three-point quadratic readout emits the exact running quark sextet, and the exact forward construction emits explicit exact forward Yukawas Y_u and Y_d. The same sextet is also realized by two supporting exact surfaces: the exact same-family witness on `current_family_only` and the restricted common-refinement transport-frame theorem on `current_family_common_refinement_transport_frame_only`. This is selected-class closure only. It does not classify all public quark frame classes.

8.1 What quarks are in this derivation

Quarks are the color-charged elementary constituents of hadronic matter. In nature, the up quark and down quark dominate protons and neutrons, while the strange quark, charm quark, bottom quark, and top quark appear in progressively heavier and more unstable sectors. In the OPH particle derivation, quarks are the only matter family whose flavor lane carries an exact theorem-grade running-mass readout together with explicit exact forward Yukawas on a public selected class.

The quark route starts from the shared flavor excitation dictionary, internalizes the target-free mass bridge on the D12 mass ray, descends the exact physical sigma datum to the selected public class f_P, applies the affine mean law, and closes both the ordered three-point readout and the exact forward construction. The same-label left-handed selector surface closes to the singleton σ_{ref}. That lower object is a negative sheet-selector statement on the selected D12 sheet. It does not replace the selected-class public theorem.

8.2 Emitted quark rows

These comparisons are made against running quark references, not direct free-particle pole masses, because quarks are confined. That is not a technicality; it is part of what the quark chapter must explain to technical readers. The exact sextet matches the official PDG 2025 API running-quark target surface, with the top coordinate taken from the PDG cross-section entry. The auxiliary direct-top row is a compare-only extraction codomain with a corpus-limited no-go boundary. The Higgs/top branch also carries the same exact top value 172.3523553288312 GeV on the declared D11 split theorem surface. The table above records the quark-family running-mass theorem surface on f_P.

8.3 Target-free mass bridge and selected-class public theorem

The same-label left-handed solver surface closes to the singleton σ_{ref} (`sigma_ref`). This is a negative same-sheet selector statement: same-sheet rephasing preserves CKM moduli, so it does not move that selected sheet to the physical CKM shell. Same-sheet overlap scans, chirality-swapped basis diagnostics, and other non-sector-attached orbit improvements are compare-only and do not change the selected-class theorem.

A separate target-free mass bridge is internalized on the emitted D12 ray. On the minimal light branch

$$y_u = c_u\,\varepsilon^6, \qquad y_d = c_d\,\varepsilon^6, \qquad \varepsilon = \frac{1}{6},$$

the light-quark overlap-defect theorem emits

$$\Delta_{ud}^{\text{overlap}} = \frac{1}{6}\log\frac{c_d}{c_u}.$$

The emitted same-family D12 mass object is the ray

$$D12_{ud}^{\text{mass}} \equiv \mathcal{R}_{D12}^{ud} \qquad (\text{paper id } \texttt{D12\_ud\_mass\_ray}),$$

and on that ray the same scalar is equivalently the one-scalar law

$$\Theta_{ud}^{\text{mass}} := \texttt{quark\_same\_family\_value\_law}.$$

The exact scalar identities are

$$\Delta_{ud}^{\text{overlap}} = \frac{t_1}{5}, \qquad \log\frac{c_d}{c_u} = \frac{6}{5}\,t_1, \qquad t_1 = 5\,\Delta_{ud}^{\text{overlap}} = \frac{5}{6}\log\frac{c_d}{c_u}.$$

The D12 mass-side package is therefore functorial:

$$\eta_Q^{\text{centered}} = -\frac{1-x_2^2}{27}\,t_1, \qquad \kappa_Q = -\frac{t_1}{54}, \qquad x_2 = -0.5175863354681689.$$

The odd source package is also forced:

$$\begin{aligned}
\beta^{\text{source}}_{u,\text{diag},B} &= \frac{t_1}{10}, \\
\beta^{\text{source}}_{d,\text{diag},B} &= -\frac{t_1}{10}, \\
\texttt{source_readback_u_log_per_side} &= \left(-\frac{t_1}{10},\,0,\,+\frac{t_1}{10}\right), \\
\texttt{source_readback_d_log_per_side} &= \left(+\frac{t_1}{10},\,0,\,-\frac{t_1}{10}\right),
\end{aligned}$$

and with the emitted spread pair

$$\sigma_u = 5.5905, \qquad \sigma_d = 3.3049,$$

one gets

$$\tau_u = \frac{\sigma_d}{10(\sigma_u+\sigma_d)}\,t_1, \qquad \tau_d = \frac{\sigma_u}{10(\sigma_u+\sigma_d)}\,t_1.$$

The mass bridge is not part of the selected-class theorem boundary. The builder-facing pure-B payload pair is a lower implementation object beneath the exact theorem surface, and `intrinsic_scale_law_D12` is the derived wrapper above $\Theta_{ud}^{\text{mass}}$.

Write the same-label left-handed physical carrier as

$$\Sigma_{ud}^{\text{phys}} := \Big\{(\sigma_{\text{id}}, \tau, U_{u,L}, U_{d,L}, V_{\text{CKM}}, I_{\text{CKM}}) : V_{\text{CKM}} = U_{u,L}^\dagger U_{d,L}\Big\}\,/\!\sim,$$

where

$$(U_{u,L}, U_{d,L}, V) \sim (U_{u,L}D_u,\ U_{d,L}D_d,\ D_u^\dagger V D_d)$$

for diagonal $D_u, D_d \in U(1)^3$. On the selected public quark frame class f_P, represented on the realized lane by the explicit common-refinement transport-frame class $[F_0^\dagger F_1]$, the exact $\Sigma_{ud}^{\text{phys}}$ element and the attached exact sigma datum are independent of the declared representative inside the selected bridge fiber. Therefore the exact sigma readout descends uniquely to target-free public data on f_P.

Define

$$\sigma_{\text{seed}}^{ud} := \frac{\sigma_u+\sigma_d}{2}, \qquad \eta_{ud} := \frac{\sigma_u-\sigma_d}{2},$$

$$A_{ud} := \frac{1}{2(1+\rho_{\text{ord}}-x_2^2)}, \qquad B_{ud} := \frac{1}{2\Big(1-x_2^2-\frac{x_2^2}{1+\rho_{\text{ord}}}\Big)}, \qquad \rho_{\text{ord}} = 1.294284936377706.$$

Then

$$g_u = g_{\text{ch}}\exp\!\big(-(A_{ud}\sigma_{\text{seed}}^{ud} - B_{ud}\eta_{ud})\big), \qquad g_d = g_{\text{ch}}\exp\!\big(-(A_{ud}\sigma_{\text{seed}}^{ud} + B_{ud}\eta_{ud})\big).$$

The affine mean law therefore emits (g_u, g_d) algebraically on f_P. The ordered three-point quadratic readout emits the exact running quark sextet, and the exact forward construction emits explicit exact forward Yukawas Y_u and Y_d on that selected public class. The selected-class public theorem can be written as

$$f_P \Longrightarrow (\sigma_u, \sigma_d, \sigma^{ud}_{\mathrm{seed}}, \eta_{ud})_{\mathrm{phys}} \Longrightarrow (g_u, g_d) \Longrightarrow (m_u, m_d, m_s, m_c, m_b, m_t),$$

together with explicit exact forward Yukawas Y_u and Y_d on the same selected-class surface.

Maximal theorem-emitted quark package theorem. Let

$$\begin{aligned} P := p_a & \\ & + (\text{Axioms 1–5}) \\ & + \text{Assumption 6} \\ & + \text{the emitted OPH nodes D1–D10} \\ & + \text{the listed public D12 quark objects.} \end{aligned}$$

Then the quark-side package can be written in four layers:

1. the emitted D12 mass ray

$$D12^{\mathrm{mass}}_{ud} = \mathcal{R}^{ud}_{D12} = \left\{ \lambda \left(\frac{1}{5}, -\frac{1 - x_2^2}{27} \right) : \lambda \geq 0 \right\}, \qquad \lambda = \mathrm{ray\_modulus} = t_1;$$

2. the same-label left-handed selector value

$$\sigma_{ud} = \sigma_{\mathrm{ref}},$$

   with canonical token `D12::same_label_left::reference_sheet`, which is a negative sheet-selector statement;

3. the separate target-free mass bridge

$$P \vdash \Delta^{\mathrm{overlap}}_{ud} = \frac{1}{6} \log \frac{c_d}{c_u}, \qquad P \vdash \Theta^{\mathrm{mass}}_{ud} = \texttt{quark\_same\_family\_value\_law};$$

4. the selected-class public theorem on f_P, which emits the target-free exact physical sigma datum, the affine pair (g_u, g_d), the exact running sextet, and explicit exact forward Yukawas Y_u and Y_d.

A separate exact same-family witness reproduces the same sextet

$$(u, d, s, c, b, t) = (0.00216, 0.00470, 0.0935, 1.273, 4.183, 172.3523553288311)\,\mathrm{GeV}$$

on `current_family_only`, and the restricted common-refinement transport-frame theorem reconstructs the same sextet and explicit exact forward Yukawas on `current_family_common_refinement_transport_frame_only`. Those supporting surfaces agree with the selected-class public theorem on the exact running target surface.

A separate D12 internal backread sidecar is on disk. On that continuation-only surface, feeding the emitted reference-free forward light-quark pair through the coefficient-ratio identification $c_u/c_d = y_u/y_d = m_u/m_d$ and the D12 normalization $\Delta^{\mathrm{overlap}}_{ud} = \frac{1}{6}\log(c_d/c_u)$ fixes

$$\Delta^{\mathrm{overlap}}_{ud} = 0.13431740757663183, \qquad t_1 = 0.6715870378831591,$$
$$\eta^{\mathrm{centered}}_{Q} = -0.018210067243314802, \qquad \kappa_Q = -0.012436796997836279.$$

This closes the D12 mass-side scalar package numerically on that sidecar surface. It does not replace the selected-class public theorem and it does not classify all public quark frame classes.

8.4 Quark Theorem Boundary

The exact closure route is

$$f_P \Longrightarrow (\sigma_u, \sigma_d, \sigma^{ud}_{\text{seed}}, \eta_{ud})_{\text{phys}} \Longrightarrow (g_u, g_d) \Longrightarrow (m_u, m_d, m_s, m_c, m_b, m_t),$$

with the exact forward construction emitting Y_u and Y_d on the same selected-class surface. This is selected-class closure only. It does not classify all public quark frame classes.

8.5 Strong-CP branch

The selected-class quark theorem fixes the running sextet and explicit forward Yukawas on the public quark frame class f_P. Strong CP is a separate phase-side invariant. The available corpus does not derive the bare QCD angle θ_{QCD}, does not emit the physical anomaly-invariant combination $\bar{\theta}$, and does not prove that the physical strong-CP phase vanishes.

The unfinished object is narrower than the quark mass theorem. The required bridge is a theorem-grade descent from exact quark/Yukawa phase data to the determinant-line phase contribution, together with a theorem fixing the topological-angle contribution on the realized branch.

9 Charged-Lepton Family Derivation

The charged-lepton derivation differs from the quark derivation in a precise way. It does not emit public charged masses from P. It emits an exact same-family witness, a closed common-shift no-go, the declared same-label q_e readback, a determinant-line lift on theorem-grade physical charged data, and a downstream algebraic mass readout from theorem-grade $A_{\text{ch}}(P)$. For any fixed formal source exponent vector $M^{\text{ch}}_{\bullet}$, the same-label readback defines a source-side determinant character

$$S_M = \sum_e M^{\text{ch}}_e \log q_e.$$

The theorem lane does not emit a theorem-grade sector-isolated charged determinant exponent vector, and it does not identify a source-side determinant character with the physical charged determinant line. The electron, muon, and tau rows are therefore reported as n/a on the public theorem lane.

The mathematical split is equally precise. The charged-lepton derivation starts from the ordered charged package, proves that the realized support is a one-dimensional linear subray, exposes the canonical quadratic support-extension direction, maps that into the charged excitation gaps, closes a two-scalar support-extension law shell, isolates the smaller eta source-readback primitive on that same carrier, and then builds the log-spectrum and forward shape/scale surface. Those centered objects are common-shift invariant. The determinant line fixes the physical affine scalar once theorem-grade physical charged data are present, and the mass readout is then algebraic. The charged theorem boundary has two distinct pieces. One piece is promotion of the latent charged sector-response candidate to theorem-grade $\widehat{C}_e$. The other is source-to-physical determinant attachment. For a fixed formal source exponent vector $M^{\text{ch}}_{\bullet}$, that attachment is the identity

$$3\mu(r) = \sum_e M^{\text{ch}}_e \log q_e(r),$$

equivalently zero determinant-normalization defect

$$N_{\text{det}}(P) = s_{\text{det}}(P) - \sum_e M^{\text{ch}}_e \log q_e(P),$$

on the charged determinant channel.

Physically, the electron is the light charged lepton that makes atoms and chemistry possible, the muon is its heavier unstable cousin, and the tau lepton is the heaviest charged lepton. The role of the charged-lepton chapter is to explain how those three rows emerge from the shared flavor dictionary without Koide-assisted fitting. The theorem-grade lane has to derive promotion of the latent charged sector-response candidate $\widehat{C}_e^{\mathrm{cand}}$ to theorem-grade $\widehat{C}_e$ by closing the branch-generator splitting theorem. On theorem-grade physical Y_e, a refinement-stable uncentered lift collapses the determinant-line section and affine anchor to one descended physical affine scalar $\mu_{\mathrm{phys}}(Y_e)$, with

$$\widetilde{C}_e(Y_e) = \widehat{C}_e(Y_e) + \mu_{\mathrm{phys}}(Y_e)\,\mathbf{1}, \qquad s_{\mathrm{det}}(Y_e) = 3\mu_{\mathrm{phys}}(Y_e), \qquad A_{\mathrm{ch}}(Y_e) = \mu_{\mathrm{phys}}(Y_e).$$

Theorem 9.1 (Charged same-carrier source-pair readback)**.** *Fix the ordered charged carrier*

$$(-1, x_2, 1), \qquad x_2 = -0.5175863354681689.$$

If the charged source pair

$$(\eta_{\mathrm{ext}}, \sigma_{\mathrm{ext}}) = (\eta_{\mathrm{source\_support\_extension\_log\_per\_side}}, \sigma_{\mathrm{source\_support\_extension\_total\_log\_per\_side}})$$

is emitted on that carrier, then the centered charged logs are

$$e_{\mathrm{log,centered}} = -\frac{(3 + x_2)\sigma_{\mathrm{ext}} - \eta_{\mathrm{ext}}}{6},$$

$$\mu_{\mathrm{log,centered}} = \frac{x_2\sigma_{\mathrm{ext}} - \eta_{\mathrm{ext}}}{3},$$

$$\tau_{\mathrm{log,centered}} = \frac{(3 - x_2)\sigma_{\mathrm{ext}} + \eta_{\mathrm{ext}}}{6},$$

and therefore the charged masses are

$$m_e = g_e\, e^{e_{\mathrm{log,centered}}}, \qquad m_\mu = g_e\, e^{\mu_{\mathrm{log,centered}}}, \qquad m_\tau = g_e\, e^{\tau_{\mathrm{log,centered}}}.$$

This theorem is exact on the same-carrier shell, but the absolute values are not emitted. The visible scalar order is

$$\eta_{\mathrm{ext}} \quad \text{then} \quad \sigma_{\mathrm{ext}},$$

and the charged absolute scale g_e is unresolved. Builder-facing code therefore exposes η_{ext} and then σ_{ext} as the first same-carrier residuals, but the theorem-grade boundary lies above that surface. If the latent candidate $\widehat{C}_e^{\mathrm{cand}}$ is promoted, then η_{ext} and σ_{ext} become charged spectral invariants instead of separate primitive goals, and the absolute-scale burden is pushed to one affine-covariant absolute charged anchor A_{ch}. In the local chain, $\widehat{C}_e$ itself is undeclared: only the centered compressed generation-bundle branch-operator candidate $\widehat{C}_e^{\mathrm{cand}}$ is on disk, and the operator-side gate for that package is the upstream promotion theorem `oph_generation_bundle_branch_generator_splitting`, not a new ad hoc charged operator ansatz. Its exact gate clause is the compression-descendant commutator statement `compression_descendant_commutator_vanishes_or_is_uniformly_quadratic_small_after_central_split`. The centered common-shift quotient is closed negatively, so centered data alone do not emit the affine anchor A_{ch}.

Theorem 9.2 (Charged absolute-scale underdetermination)**.** *Let*

$$E_e^{\mathrm{centered}} = (e_{\mathrm{log,centered}}, \mu_{\mathrm{log,centered}}, \tau_{\mathrm{log,centered}})$$

be the centered charged log triple emitted from the charged source pair $(\eta_{\rm ext}, \sigma_{\rm ext})$*. Then for every* $c \in \mathbb{R}$*,*

$$Y_e(c) := \exp(c)\,\mathrm{diag}(e^{E_e^{\rm centered}})$$

has the same charged spectral invariants

$$\eta_{\rm ext}, \qquad \sigma_{\rm ext}, \qquad \gamma_{21}, \qquad \gamma_{32},$$

the same centered log vector, and the same ratio data. Only the absolute masses scale:

$$(m_e, m_\mu, m_\tau) \longmapsto e^c\,(m_e, m_\mu, m_\tau).$$

Hence the charged OPH chain determines only the quotient class

$$E_e^{\rm centered} \in \mathbb{R}^3/\langle(1,1,1)\rangle,$$

not the absolute scalar $g_e = e^c$*.*

The proof is immediate from the centered sum rule

$$e_{\rm log,centered} + \mu_{\rm log,centered} + \tau_{\rm log,centered} = 0,$$

which implies

$$\det Y_e^{\rm shape} = 1, \qquad \det Y_e = g_e^3.$$

All emitted charged invariants depend only on differences of log entries, so a common shift in the $(1,1,1)$ direction leaves them unchanged. This is the exact reason the available corpus closes the P-driven charged lane as a no-go rather than as a mass theorem: the theorem surface fixes the centered shape, while the determinant-line landing fixes the absolute normalization.

The charged absolute-scale lane is explicitly typed. The charged scale is a linear quantity

$$g_e = e^{\mu_e^{\rm abs}},$$

so the log-coordinate $\mu_e^{\rm abs}$ must not be mixed directly with centered log gaps. The same-family writeback therefore records the type-consistent shell

$$\mu_{e,\rm seed}^{\rm abs} = \log(0.9231656602589082) = -0.07994658034676537,$$

$$\mu_{e,\rm cand}^{\rm abs} = \mu_{e,\rm seed}^{\rm abs} - \gamma_{\min} = -0.38231224060567365, \qquad g_e^{\rm cand} = e^{\mu_{e,\rm cand}^{\rm abs}} = 0.6822819838027987.$$

This is a representation-consistency shell only, not a charged-mass theorem. It fixes the linear-vs-log coordinate discipline. The emitted value law for g_e is work in progress. The audited shortcut

$$\Delta_e^{\rm abs} = 0.30236566025890826, \qquad g_e = 0.6822819838027987$$

is not a theorem-grade closure. It merely chooses one representative on the common-shift orbit and does not land on the physical charged masses.

The next theorem object beyond this shell is one theorem-grade affine-covariant section $A_{\rm ch}$ of the quotient map, satisfying

$$A_{\rm ch}(\ell + c\,\mathbf{1}) = A_{\rm ch}(\ell) + c.$$

The clean realization of that section is an uncentered charged response lift carrying a determinant line, in which

$$A_{\rm ch} = \tfrac{1}{3}\log\det(Y_e) = \tfrac{1}{3}\,\mathrm{tr}(\log Y_e).$$

Once such an anchor exists,

$$g_e = e^{A_{\rm ch}(\ell)}, \qquad \Delta_e^{\rm abs} = \log g_{\rm ch}^{\rm shared} - A_{\rm ch}(\ell).$$

Compare-only charged continuation bridge. For comparison, there is a sharper charged continuation bridge under three extra continuation assumptions:

1. a uniform $\mathbb{Z}_6$ center-label ensemble, so $\varepsilon = 1/6$;
2. the balanced Hermitian circulant branch, so $|b|/a = 1/\sqrt{2}$;
3. the OPH phase choice $\delta = 2/9$.

With the scale-free normalization $a = 1$, the ordered roots

$$r_k = a + 2|b| \cos\left(\delta + \frac{2\pi k}{3}\right)$$

become

$$r_e = 0.040349908219207475, \qquad r_\mu = 0.5802119201475368, \qquad r_\tau = 2.3794381716332555.$$

Under those assumptions the bridge selects the same-carrier pair

$$\eta_{\text{ext}} = -6.729586682888832, \qquad \sigma_{\text{ext}} = 8.154061112725994,$$

with ordered gap values

$$\gamma_{21} = 5.33160859254774, \qquad \gamma_{32} = 2.822452520178255, \qquad \kappa_{\text{ext}} = -4.59605680397234,$$

and centered logs

$$E_{\text{log,centered}} = [-4.495223235091244,\ 0.836385357456495,\ 3.6588378776347503].$$

Against the charged references, this continuation-centered shape has residual norm

$$\left\| E^{\text{cont}}_{\text{log,centered}} - E^{\text{ref}}_{\text{log,centered}} \right\| \approx 2.13 \times 10^{-5},$$

so this compare-only branch is a near-exact centered-shape closure up to one common absolute scale. It does not promote the theorem lane: the centered charged shape is nearly exact on this branch, but the public affine absolute anchor is external. The compare-only common scale required for exact absolute masses is

$$g_e^\star = 0.04577885783568762.$$

Equivalently, relative to the stored shared-budget seed

$$g_{\text{ch}}^{\text{shared}} = 0.9231656602589082,$$

the missing affine absolute anchor on the determinant-line route would have to take the target value

$$\Delta_e^{\text{abs},\star} = \log \frac{g_{\text{ch}}^{\text{shared}}}{g_e^\star} = 3.003986333402356.$$

This identifies the public charged theorem boundary: the compare-only branch nearly solves the centered charged shape, but the theorem-grade derivation lacks promotion of the latent candidate $\widehat{C}_e^{\text{cand}}$ to theorem-grade $\widehat{C}_e$, and then one affine-covariant absolute anchor A_{ch} that would turn that centered readback into public charged masses on the theorem lane.

10 Neutrino Family Derivation

The neutrino derivation is the strongest positive flavor-side closure after the Higgs/top critical stage. The intrinsic isotropic branch is excluded exactly. On the declared weighted-cycle theorem branch, the PMNS/hierarchy shape lands in the observed regime, the bridge-rigidity theorem emits C_ν, absolute attachment fixes the paper-facing amplitude B_ν and λ_ν, and that branch emits the absolute eV-scale family. The legacy one- and two-parameter adapters are compare-only diagnostic continuations beneath that theorem branch.

On that weighted-cycle theorem branch the paper emits

$$\theta_{12} = 34.225904631810025^\circ, \qquad \theta_{23} = 49.72282845058266^\circ, \qquad \theta_{13} = 8.686355527700156^\circ,$$

$$\delta_{\rm PMNS} = 305.58061231449796^\circ, \qquad J = -0.02753115613565372,$$

and the dimensionless hierarchy invariant

$$\frac{\Delta m^2_{21}}{\Delta m^2_{32}} = 0.030721110097966534.$$

Once the repaired weighted-cycle Majorana matrix is transported explicitly through the closed shared same-label basis by

$$M_{\rm shared} = U^*_{e,{\rm left}}\, M_{\rm wc}\, U^\dagger_{e,{\rm left}}, \qquad U_{\nu,{\rm shared}} = U_{e,{\rm left}}\, U_{\rm wc},$$

the physical PMNS path is recovered on that anchored branch as

$$U_{\rm PMNS} = U^\dagger_{e,{\rm left}} U_{\nu,{\rm shared}} = U_{\rm wc}.$$

Canonical Takagi congruence of that transported symmetric matrix, with

$$U^T_{\rm PMNS} M_{\rm shared} U_{\rm PMNS} = {\rm diag}(m_i) \in \mathbb{R}_{>0}$$

and the electron-row gauge $U_{e1} \in \mathbb{R}_{>0}$, therefore emits one physical Majorana pair on the declared weighted-cycle transport branch:

$$\alpha^{\rm (Maj)}_{21} = 153.6185177794357^\circ, \qquad \alpha^{\rm (Maj)}_{31} = 257.0032408220805^\circ.$$

The exponent law is fixed by the positive transport-load segment between $\chi = 1 + \epsilon$ and $1 + \gamma_{1/2}$: on a one-dimensional affine segment, the balanced selector and the least-distortion selector for any positive translation-invariant quadratic form coincide at the midpoint, so

$$D_\nu = \frac{\chi + (1 + \gamma_{1/2})}{2}, \qquad p = 1 + \gamma + \frac{\epsilon}{D_\nu}.$$

On the emitted positive selector segment there is a stronger compare-only two-parameter adapter: solving τ_ν against the representative central ratio and then solving λ_ν against Δm^2_{32} gives

$$(m_1, m_2, m_3) = (0.01745663295,\ 0.01948419960,\ 0.05308139066)\ {\rm eV},$$

$$\Delta m^2_{21} = 7.49 \times 10^{-5}\ {\rm eV}^2,$$
$$\Delta m^2_{31} = 2.5129 \times 10^{-3}\ {\rm eV}^2,$$
$$\Delta m^2_{32} = 2.438 \times 10^{-3}\ {\rm eV}^2.$$

These exact central numbers are compare-only because the proof-facing theorem lane is the weighted-cycle bridge-rigid branch with

$$C_\nu = 0.9994295999075177, \qquad P_\nu = 6.699825740519345, \qquad B_\nu = P_\nu C_\nu = 6.696004159297337,$$

and therefore

$$\lambda_\nu = \frac{m_{\star,\mathrm{eV}}}{q_{\mathrm{mean}}^{p_\nu}} P_\nu C_\nu = 1.7237014208357415,$$

$$m_i = \lambda_\nu \widehat{m}_i, \qquad \Delta m_{ij}^2 = \lambda_\nu^2 \widehat{\Delta}_{ij}.$$

The normalized same-label overlap-defect weight section $qbar_e$ is closed beneath this step, and on the isotropic branch the same-label phase-cocycle theorem, the selector-centered common-refinement bundle descent, and the finite-angle centered edge-norm theorem are also closed, so the exact two-parameter positive-segment adapter, the bridge corridor, and the bridge-coordinate sidecars are diagnostic surfaces beneath the bridge-rigidity invariant C_ν and its induced paper-facing amplitude B_ν rather than proof-facing ingredients.

$$C_\nu := \frac{B_\nu}{I_\nu^{1/2}\, \widehat{\mathrm{ratio}}^{\,1/2}\, \mathrm{sum\_defect}^{-1}}$$

above the emitted proxy and closed normalizer. The attached stack also factors exactly through $q_e = q_{\mathrm{mean}} qbar_e$, so a $qbar_e$-only collapse law for the bridge factor does not follow from the theorem surface stated here. The best closed constructive local object beneath that bridge is the defect-weighted same-label edge family $q_e = \sqrt{g_e d_e}$ together with the induced μ_e family, but that family sits below C_ν and the induced paper-facing amplitude B_ν rather than replacing them. The weighted-cycle branch fixes the PMNS data, the hierarchy ratio, and the scale-free mass normal form; on that branch the bridge-rigidity theorem fixes C_ν above the emitted proxy P_ν, and absolute attachment then emits the absolute eV-scale family.

Cosmological-neutrino pressure surface. The same weighted-cycle absolute-attachment branch implies

$$\sum_i m_{\nu_i} = 0.09001192964464505 \text{ eV}.$$

Under standard relic-neutrino inheritance, and absent an explicitly declared extra OPH relativistic coherence channel, the cosmological-neutrino-background branch uses

$$N_{\mathrm{eff}}^{\mathrm{OPH}} = 3.044,$$

with the usual broad, low-amplitude normal-ordering free-streaming imprint. This is a cosmological pressure surface for the weighted-cycle branch, not an input to its theorem proof. The emitted mass sum sits below the Planck+BAO 0.12 eV bound [6] but is exposed to strict DESI DR2 ΛCDM-style bounds near 0.0642 eV under their model assumptions [7]. A robust cosmological upper bound below the emitted value, in a model class that also respects oscillation lower limits and the OPH background assumptions, would directly pressure the weighted-cycle absolute-mass branch.

On the same-label neutrino-only branch, the centered eta-class is exactly S_3-isotropic, so the same-label data are edge-constant and the solar $1-2$ split cannot open there by itself. The first solar mover therefore has to come from realized flavor-side same-label gap/defect readback, not from another neutrino-only selector tweak.

10.1 Intrinsic neutrino eta-chain

The proof-facing input is smaller than the raw same-label matrix payload. It compresses to the same-label scalar certificate

$$(g_e,\ \omega_e)_{e\in\{12,23,31\}}, \qquad \omega_e = \text{same-label overlap}_e^2, \qquad d_e = 1-\omega_e,$$

modulo one common scale. From that certificate one forms

$$q_e = \sqrt{g_e d_e}, \qquad \eta_e = \log q_e - \frac{1}{3}\sum_f \log q_f, \qquad e \in \{12,23,31\},$$

the intrinsic selector depends only on the centered class

$$[\eta_e] \in \mathbb{R}^3/\langle(1,1,1)\rangle.$$

Equivalently one may work with any positive normalized family

$$\mu_e = \frac{e^{\eta_e}}{\frac{1}{3}\sum_f e^{\eta_f}}.$$

Common scaling cancels identically, so the intrinsic selector is determined by the centered eta-class alone.

On the isotropic neutrino-only branch one has

$$(\eta_{12}, \eta_{23}, \eta_{31}) = (0,0,0),$$

which emits the intrinsic masses

$$(m_1, m_2, m_3) = (2.3986448447627196,\ 2.3986448447627196,\ 2.590074050773907) \times 10^{-12}\ \text{GeV},$$

with

$$\Delta m_{21}^2 = 0, \qquad \Delta m_{31}^2 = 9.549864971855843 \times 10^{-25}\ \text{GeV}^2.$$

This is the exact neutrino-only isotropy obstruction behind the solar-splitting boundary.

Theorem 10.1 (Same-label scalar certificate suffices for intrinsic neutrino masses)**.** *Assume the same-label gap and overlap scalars are emitted on all realized arrows. Then the full intrinsic neutrino mass-eigenstate bundle factors through the scalar certificate*

$$(g_e, \omega_e)_{e\in\{12,23,31\}}$$

or equivalently through the centered eta-class $[\eta_e]$. No additional raw matrix payload is needed to form the intrinsic selector, the depressed-cubic spectrum, the ordered splittings, or the intrinsic neutrino-side basis.

Theorem 10.2 (Exact principal-branch selector from the centered eta-class)**.** *Fix the cycle sum*

$$\Omega = \psi_{12} + \psi_{23} + \psi_{31}, \qquad |\Omega| < \frac{3\pi}{2},$$

and a positive weight family μ_e. On the principal branch $\psi_e \in (-\pi/2, \pi/2)$, the affine energy

$$A(\psi) = \sum_e \mu_e (1-\cos\psi_e)$$

has a unique minimizer. It is given by

$$\psi_e^\star = \arcsin(\lambda/\mu_e),$$

where λ is the unique solution of

$$\sum_e \arcsin(\lambda/\mu_e) = \Omega, \qquad \lambda \in (-\min_e \mu_e, \min_e \mu_e).$$

Proof. The Euler–Lagrange equations are $\mu_e \sin\psi_e = \lambda$. On the principal branch the scalar function

$$F(\lambda) = \sum_e \arcsin(\lambda/\mu_e)$$

is strictly increasing because

$$F'(\lambda) = \sum_e \frac{1}{\sqrt{\mu_e^2 - \lambda^2}} > 0.$$

Its range is $(-3\pi/2, 3\pi/2)$, so $F(\lambda) = \Omega$ has a unique solution. Strict convexity follows because the Hessian of A is

$$\nabla^2 A = \mathrm{diag}(\mu_e \cos\psi_e),$$

which is positive definite on the principal branch and is positive definite after restriction to the affine plane $\sum_e \psi_e = \Omega$. □

Let

$$a = m_\star = \frac{v^2}{\mu_u}, \qquad \rho = |(M_0)_{12}|.$$

Then the intrinsic Majorana matrix is

$$M(\eta) = \begin{pmatrix} a & \rho e^{i\psi_{12}} & \rho e^{i\psi_{31}} \\ \rho e^{i\psi_{12}} & a & \rho e^{i\psi_{23}} \\ \rho e^{i\psi_{31}} & \rho e^{i\psi_{23}} & a \end{pmatrix}.$$

Theorem 10.3 (Exact intrinsic spectral cubic)**.** *Let $H = M^\dagger M$. Then*

$$H = dI + T, \qquad d = a^2 + 2\rho^2,$$

where T has zero diagonal and off-diagonals

$$x_{12} = 2a\rho\cos\psi_{12} + \rho^2 e^{i(\psi_{23}-\psi_{31})},$$

$$x_{23} = 2a\rho\cos\psi_{23} + \rho^2 e^{i(\psi_{31}-\psi_{12})},$$

$$x_{13} = 2a\rho\cos\psi_{31} + \rho^2 e^{i(\psi_{12}-\psi_{23})}.$$

Define

$$P = |x_{12}|^2 + |x_{23}|^2 + |x_{13}|^2, \qquad Q = \Re(x_{12}x_{23}\overline{x_{13}}).$$

Then the eigenvalues λ_k of T are exactly the three real roots of

$$\lambda^3 - P\lambda - 2Q = 0,$$

and the intrinsic squared masses are

$$m_k^2 = d + \lambda_k.$$

Equivalently,

$$\lambda_k = 2\sqrt{\frac{P}{3}}\cos\left(\frac{1}{3}\arccos\left(\frac{3\sqrt{3}\,Q}{P^{3/2}}\right) - \frac{2\pi k}{3}\right).$$

Corollary 10.4 (Intrinsic eta-chain closure)**.** *Once the centered eta-class is emitted at the flavor boundary, the intrinsic neutrino branch emits*

$$\nu_1, \nu_2, \nu_3, \qquad \Delta m^2_{21}, \Delta m^2_{31}, \Delta m^2_{32},$$

the ordering, and the intrinsic neutrino-side basis U_ν with no PMNS import and no flavor-label leakage.

10.2 Perturbative laws around the isotropic point

Write

$$\psi_e = \varphi + \delta_e, \qquad \delta_{12} + \delta_{23} + \delta_{31} = 0.$$

At first order in the centered eta-class,

$$\delta_e = -\tan\varphi\, \eta_e + O(\eta^2).$$

Define

$$\sigma^2 = \frac{2}{3}\sum_e \delta_e^2.$$

Then

$$m^2_{1,2} = m_d^2 \mp 2a\rho \sin\varphi\, \sigma + O(\delta^2),$$

so

$$\Delta m^2_{21} = 4a\rho \sin\varphi\, \sigma + O(\delta^2) = 4a\rho \frac{\sin^2\varphi}{\cos\varphi}\sqrt{\frac{2}{3}\sum_e \eta_e^2} + O(\eta^2).$$

The ordered atmospheric gaps obey the corrected first-order laws

$$\Delta m^2_{31} = \Delta_{\mathrm{atm,iso}} + \frac{1}{2}\Delta m^2_{21} + O(\delta^2), \qquad \Delta m^2_{32} = \Delta_{\mathrm{atm,iso}} - \frac{1}{2}\Delta m^2_{21} + O(\delta^2),$$

so the first-order invariant atmospheric object is the collective-to-doublet centroid gap

$$\Delta_{\mathrm{cent}} = m_3^2 - \frac{1}{2}(m_1^2 + m_2^2) = \Delta_{\mathrm{atm,iso}} + O(\delta^2).$$

Its quadratic shift is

$$\delta\Delta_{\mathrm{cent}} = -a\rho\,\sigma^2 \frac{a(4\cos^2\varphi - 1) + 6\rho\cos\varphi}{2a\cos\varphi + \rho} + O(\delta^3).$$

The largest-mass singular vector obeys the first-order deformation law

$$u_3 = u + \kappa \begin{pmatrix} \delta_{23} \\ \delta_{31} \\ \delta_{12} \end{pmatrix} + O(\delta^2), \qquad u = \frac{1}{\sqrt{3}}(1,1,1)^T, \qquad \kappa = \frac{\sqrt{3}\,(2a\sin\varphi + 3i\rho)}{9(2a\cos\varphi + \rho)}.$$

Equivalently,

$$u_3 = u - \tan\varphi\, \kappa \begin{pmatrix} \eta_{23} \\ \eta_{31} \\ \eta_{12} \end{pmatrix} + O(\eta^2).$$

One exact demonstration uses centered eta-class

$$(\eta_{12}, \eta_{23}, \eta_{31}) = (0.13631072512014578,\ -0.18362987264261327,\ 0.0473191475224675),$$

and emits intrinsic masses

$$(m_1, m_2, m_3) = (2.3929601069646055,\ 2.4048200109774875,\ 2.589606227283229) \times 10^{-12}\ \text{GeV},$$

with

$$\Delta m^2_{21} = 5.690121167370743 \times 10^{-26}\ \text{GeV}^2, \qquad \Delta m^2_{31} = 9.798023388600230 \times 10^{-25}\ \text{GeV}^2.$$

These are exact outputs of the intrinsic eta-chain once the centered eta-class is supplied. They are not flavor-labeled OPH rows.

10.3 Weighted-cycle bridge rigidity and absolute attachment

The weighted-cycle route closes the PMNS and hierarchy steps. On that lane the bridge-rigidity theorem fixes the reduced invariant C_ν above the emitted proxy P_ν, with $B_\nu = P_\nu C_\nu$ retained only as the paper-facing amplitude parameterization:

$$\begin{aligned} C_\nu &= G^2 Q S^{-1/2} = 0.9994295999075177, \\ P_\nu &= 6.699825740519345, \\ B_\nu &= P_\nu C_\nu = 6.696004159297337. \end{aligned}$$

The absolute-attachment theorem therefore emits

$$\begin{aligned} \lambda_\nu &= \frac{m_{\star,\text{eV}}}{q_{\text{mean}}^{p_\nu}}\, P_\nu C_\nu = 1.7237014208357415, \\ m_i &= \lambda_\nu \hat{m}_i, \\ \Delta m^2_{ij} &= \lambda_\nu^2 \widehat{\Delta m^2_{ij}}. \end{aligned}$$

No hidden discrete branch is present on that lane. The one-parameter atmospheric-anchor slice and the stronger two-parameter positive-segment adapter are only diagnostic compare-only continuations beneath the bridge-rigidity invariant C_ν and its induced paper-facing amplitude B_ν.

11 Hadrons, QCD, and the Emergence of Ordinary Matter

The hadron derivation is the most operationally demanding part of the particle program. The source-only artifact is a Ward-projected spectral-measure contract for a backend. A source-only hadron mass row requires a working OPH hadron backend, such as GLORB/Echosahedron, that can emit real Ward-projected hadronic spectral data with manifest provenance and production systematics. Local surrogates, ordinary Chrome/Oracle workers, and bookkeeping scripts cannot promote this sector. The source-backend boundary has an empirical closure policy. The empirical closure row class uses a separate $e^+e^- \to$ hadrons payload class for the hadronic spectral contribution and stays separate from source-only OPH rows. The source-only derivation nevertheless keeps the mathematical execution bridge, deterministic runtime receipt, writeback/evaluation path, and surrogate validation machinery as non-promoting scaffolding for the backend lane.

11.1 Seeded $2+1$ family and unquenched measure

Let

$$m_l := \frac{m_u + m_d}{2}, \qquad \rho_l := \frac{m_u + m_d}{2\,\Lambda^{(3)}_{\overline{\rm MS}}}, \qquad \rho_s := \frac{m_s}{\Lambda^{(3)}_{\overline{\rm MS}}}.$$

The seeded fixed-physics family is

$$a\Lambda_n = a\Lambda_{\rm seed}\, 2^{-n}, \qquad a\Lambda_{\rm seed} = \sqrt{\rho_l \rho_s},$$

$$am_l^{(n)} = a\Lambda_n \rho_l, \qquad am_s^{(n)} = a\Lambda_n \rho_s,$$

$$\beta_n = 6 + \frac{9}{2\pi^2} \log \frac{1}{a\Lambda_n}, \qquad L_n = \left\lceil \frac{\lambda_L^{\rm target}}{a\Lambda_n} \right\rceil, \qquad T_n = \left\lceil \frac{\lambda_T^{\rm target}}{a\Lambda_n} \right\rceil.$$

The unquenched ensemble measure is

$$d\mu_n(U) = Z_n^{-1} \exp[-S_g(U;\beta_n)] \det D_l(U; am_l^{(n)})^2 \det D_s(U; am_s^{(n)})\, dU.$$

On this branch, $N_f = 2+1$ and QED is off.

11.2 Runtime receipt and deterministic cfg/source contract

The emitted execution law is

$$U_{n,c} = K_n^{N_{\rm therm} + ({\rm cfg\_index})\, N_{\rm sep}}(U_{\rm cold}; {\rm seed}_{n,c}),$$

with stop-time formula

$$t_{\rm stop}(n,c) = N_{\rm therm} + ({\rm cfg\_index})\, N_{\rm sep}.$$

The deterministic cfg seed law is

$${\rm seed}_{n,c} = {\rm bytes.fromhex}({\rm cfg\_seed\_hash}_{n,c}),$$

with

$${\rm cfg\_seed\_hash} = {\rm SHA256}\Big({\rm Serialize}\big({\rm ensemble\_id}, \beta, L, T, am_l, am_s, {\rm cfg\_index}\big)\Big).$$

The source set is fixed as

$$S_{n,c} = \{[0,0,0,0],\ [L_n//2, L_n//2, L_n//2, T_n//2]\}.$$

The non-null external runtime receipt used in the surrogate execution is

$$N_{\rm therm} = 2048, \qquad N_{\rm sep} = 512.$$

These values are surrogate execution inputs only; they are not claimed as theorem outputs.

The production geometry summary makes the runtime boundary concrete. On the emitted seeded $2+1$ family there are three ensembles and six configurations total. If one stores all four links at every site as full double-complex 3×3 matrices, the naive raw gauge storage estimate is

$$576 \text{ bytes/site},$$

which gives total naive raw gauge storage

$$2.80071464105088 \times 10^{14} \text{ bytes}$$

over the production schedule. The normalized backend correlator dump needed by the analysis layer is tiny:

$$195264 \text{ bytes}$$

for the full $\pi_{\rm iso}$, $N_{\rm iso,dir}$, and $N_{\rm iso,ex}$ payload on the production schedule. The hadron requirement is the backend export bundle that would feed the normalized dump from real production execution.

11.3 Operational reading of the required computation

It is useful to separate the lightweight analysis layer from the missing backend layer. Informally, the repository-side control plane is explicit: the seeded ensemble family, the deterministic cfg/source contract, the export manifest, the jackknife evaluator, the forward-window selector, and the publication-budget schema are all explicit. The missing step is the production export of Ward-projected hadronic spectral data from a real OPH hadron backend.

Technically, the required production computation is the standard lattice-QCD stable-channel workflow on the emitted family: run unquenched RHMC/HMC for the three seeded ensembles; for each realized cfg and fixed source solve the clover-improved Wilson light/strange systems; construct the zero-momentum π_{iso}, $N_{\text{iso,dir}}$, and $N_{\text{iso,ex}}$ two-point sequences; write the backend bundle; convert that bundle into the repo-side production dump; then let the existing jackknife and forward-window machinery emit $am_{X,\text{ground}}$, R_X, $m_X[\text{GeV}]$, and the published σ_{stat}, δ_{cont}, δ_{vol}, and δ_χ fields. This gives a complete source-only execution contract and a separate empirical display policy.

The engineering burden is therefore asymmetric. Local execution is sufficient for manifest generation, surrogate validation, and downstream readout tests, because those stages operate on a tiny correlator payload. The physical branch is different: it needs backend hardware and execution semantics that are absent from the repository environment. This paper therefore treats GLORB/Echosahedron-class OPH hardware, or an equivalent working OPH hadron backend, as the source-only promotion gate. Production hadron masses enter the particle pipeline only with a backend-emitted Ward-projected spectral measure and its uncertainty budget. This is a source-backend scope boundary.

11.4 Stable-channel correlators, effective masses, and forward windows

The pion stable channel is

$$p_\pi^{(n,c,s)}(t) = \sum_x \Re\,\text{tr}_{c,\text{spin}}\left[\gamma_5 S_l(x;s)\gamma_5 S_l(s;x)\right].$$

The nucleon stable channel is

$$p_{N,\text{dir}}^{(n,c,s)}(t) = \sum_x G_d, \qquad p_{N,\text{ex}}^{(n,c,s)}(t) = \sum_x G_x,$$

$$p_N^{(n,c,s)}(t) = p_{N,\text{dir}}^{(n,c,s)}(t) - p_{N,\text{ex}}^{(n,c,s)}(t).$$

Cfg/source and ensemble averaging are

$$\bar{p}_X^{(n,c)}(t) = \frac{1}{|S_{n,c}|}\sum_{s\in S_{n,c}} p_X^{(n,c,s)}(t),$$

$$C_X^{(n)}(t) = \frac{1}{|C_n|}\sum_{c\in C_n} \bar{p}_X^{(n,c)}(t).$$

The effective-mass laws are

$$am_{\text{eff},\pi}(t) = \log\frac{C_\pi(t)}{C_\pi(t+1)}, \qquad am_{\text{eff},N}(t) = \log\frac{|C_N(t)|}{|C_N(t+1)|}.$$

The forward window is

$$W_n = \{t : 1 \le t+1 < \lfloor T_n/2 \rfloor\}.$$

The evaluator monitors log-convexity,

$$R_{\log\mathrm{conv},\pi}(t) = C_\pi(t)^2 - C_\pi(t-1)C_\pi(t+1),$$

$$R_{\log\mathrm{conv},N}(t) = |C_N(t)|^2 - |C_N(t-1)|\,|C_N(t+1)|,$$

tail-drop,

$$D_X(t) = am_{\mathrm{eff},X}(t) - am_{\mathrm{eff},X}(t+1),$$

and mirror suppression,

$$M_X(t) = \exp[-am_{\mathrm{eff},X}(t)(T_n - 2t)].$$

The selected forward window is the longest contiguous run satisfying finite effective masses, non-negative log-convexity up to tolerance, nonnegative tail-drop up to tolerance, mirror suppression below threshold, and local plateau flatness. The candidate ground-state mass is then the weighted window average

$$am_{X,\mathrm{ground}} = \frac{\sum_{t\in W_X^{\mathrm{sel}}} w_t\, am_{\mathrm{eff},X}(t)}{\sum_{t\in W_X^{\mathrm{sel}}} w_t}, \qquad w_t = \frac{1}{\max(\sigma_t^2, \varepsilon)}.$$

11.5 Statistics, systematics, and dimensional readout

Delete-1 jackknife is performed over the cfg axis after source averaging inside each cfg. With n_{cfg} configurations and integrated autocorrelation time $\tau_{\mathrm{int,cfg}}$, the effective cfg count is

$$n_{\mathrm{eff,cfg}} = \frac{n_{\mathrm{cfg}}}{2\,\tau_{\mathrm{int,cfg}}}.$$

The published statistical error is

$$\sigma_{\mathrm{stat},X} = \mathrm{JKstderr}(am_{X,\mathrm{ground}}).$$

The machine-readable systematics field uses

$$\sigma_{\mathrm{sys},X} = \sqrt{\delta^2_{\mathrm{cont},X} + \delta^2_{\mathrm{vol},X} + \delta^2_{\chi,X}}.$$

Continuum, volume, and chiral proxies are encoded as

$$R_X^{(n)} = \frac{am_X^{(n)}}{a\Lambda_n}, \qquad R_X^{(n)} \approx R_X(0) + c_X(a\Lambda_n)^2,$$

$$\delta^{(n)}_{\mathrm{cont},X} = a\Lambda_n \left| R_X^{(n)} - R_X(0) \right|,$$

$$\delta^{(n)}_{\mathrm{vol},\pi} = am_\pi^{(n)} \frac{e^{-am_\pi^{(n)} L_n}}{\max(am_\pi^{(n)} L_n, 1)},$$

$$\delta^{(n)}_{\mathrm{vol},N} = am_N^{(n)} \frac{e^{-am_\pi^{(n)} L_n}}{\max(am_\pi^{(n)} L_n, 1)},$$

$$\delta^{(n)}_{\chi,\pi} = a\Lambda_n \left| R_\pi^{(n)} - \langle R_\pi \rangle_n \right|,$$

$$Q_N^{(n)} = \frac{R_N^{(n)}}{R_\pi^{(n)}}, \qquad \delta^{(n)}_{\chi,N} = a\Lambda_n \langle R_\pi \rangle_n \left| Q_N^{(n)} - \langle Q_N \rangle_n \right|.$$

These are surrogate publication proxies, not production physical systematics. Given a ground-state candidate,

$$R_X = \frac{am_{X,\text{ground}}}{a\Lambda^{(3)}_{\overline{\text{MS}}}}, \qquad m_X[\text{GeV}] = R_X\,\Lambda^{(3)}_{\overline{\text{MS}}}[\text{GeV}],$$

with

$$\Lambda^{(3)}_{\overline{\text{MS}}} = 0.3344017073\ \text{GeV}.$$

11.6 Surrogate execution bundle and frontier

The surrogate execution evolves a latent state $z \in \mathbb{R}^d$ with Hamiltonian

$$H(z,p) = S(z) + \frac{1}{2}\sum_i p_i^2,$$

where

$$S(z) = \frac{1}{2}\sum_i \omega_i^2 z_i^2 + \lambda_4 \sum_i z_i^4 + \kappa \sum_i (z_{i+1} - z_i)^2 + 2\alpha_l \sum_i \log(\mu_l^2 + z_i^2) + \alpha_s \sum_i \log(\mu_s^2 + \tfrac{1}{2} z_i^2).$$

Leapfrog integration plus Metropolis accept/reject gives the surrogate HMC update

$$(z,p) \mapsto (z',p'), \qquad P_{\text{acc}} = \min(1, e^{-\Delta H}).$$

This kernel is not physical lattice-QCD RHMC/HMC. It is a deterministic executable surrogate honoring the emitted receipt and seed law.

For validation of the execution bridge, the surrogate locks the ground-state masses to the hadron audit proxy values

$$m_{\pi,\text{proxy}} = 0.13497682776768472\ \text{GeV}, \qquad m_{N,\text{iso,proxy}} = \frac{m_p + m_n}{2} = 0.93891875434\ \text{GeV}.$$

Thus

$$am^{(n)}_{\pi,\text{proxy}} = \frac{m_{\pi,\text{proxy}}}{\Lambda^{(3)}_{\overline{\text{MS}}}[\text{GeV}]}\, a\Lambda_n,$$
$$am^{(n)}_{N,\text{iso,proxy}} = \frac{m_{N,\text{iso,proxy}}}{\Lambda^{(3)}_{\overline{\text{MS}}}[\text{GeV}]}\, a\Lambda_n.$$

The surrogate correlator is

$$C_X^{\text{sur}}(t) = A_{X,0}e^{-am_X t} + A_{X,1}e^{-am_{X,\text{ex}} t} + A_{X,\text{mir}}e^{-am_X(T_n - t)}$$

up to a multiplicative correlated noise factor. Direct and exchange nucleon pieces are written as

$$C^{\text{sur}}_{N,\text{dir}}(t) = f_{\text{dir}}\, C_N^{\text{sur}}(t), \qquad C^{\text{sur}}_{N,\text{ex}}(t) = (f_{\text{dir}} - 1)\, C_N^{\text{sur}}(t),$$

so that

$$C_N^{\text{sur}}(t) = C^{\text{sur}}_{N,\text{dir}}(t) - C^{\text{sur}}_{N,\text{ex}}(t).$$

On the finest surrogate ensemble, the resulting diagnostic candidates are

$$m^{\text{sur}}_{\pi,\text{iso}} = 0.135039383836\ \text{GeV}, \qquad m^{\text{sur}}_{N,\text{iso}} = 0.938960210578\ \text{GeV},$$

with worst absolute error approximately 6.26×10^{-5} GeV across the surrogate stable-channel family. These numbers validate the emitted execution bridge. They are not promotable production hadron predictions.

So the hadron derivation closes the full

$$\text{receipt} \to \text{execution} \to \text{writeback} \to \text{evaluation} \to \text{budgets} \to \text{forward-window summary}$$

path on executed surrogate data, while the physical closure requires production unquenched RHMC/HMC, real Dirac solves and baryon contractions, production autocorrelation studies, and production continuum / finite-volume / chiral systematics.

12 Observer-Centric Particle Ontology and Measurement

This paper cannot stop at masses and couplings. OPH is observer-centric at the level of its basic ontology, so a particle chapter also has to explain what a particle *is* in this language and how a measurement turns an excitation surface into an actual observed state. Fortunately, this is one of the places where the OPH theorem surface contains real mathematical content rather than only interpretation.

12.1 Observer patches, overlap consistency, and particle data

Ref. [2] formulates the basic kinematic picture in patch-net language: each observer patch carries a local state, neighboring patches compare a shared interface alphabet, and a global state is physically admissible exactly when neighboring projections agree on the overlap. The algebraic OPH version of the same statement is Axioms 4.1 and 4.2: local physical data are carried by patch algebras $\mathcal{A}(P)$, and those local states must agree on shared subalgebras whenever patches overlap.

This matters for particles because OPH does not begin with a single absolute global particle basis. It begins with patch-local algebras, overlap-visible observables, edge sectors, and transport data. A particle row in the reported derivation is therefore not best understood as an isolated primitive label. It is a readout claim about a stable or candidate-stable excitation structure visible to a family of observer patches and consistent on their overlaps.

In the observer formulation, an observer is represented by

$$O = (P, \mathcal{A}(P), \rho, R),$$

where P is the observer patch, $\mathcal{A}(P)$ its local algebra, ρ the local state, and R the record algebra. For the particle derivation, R belongs to the quantum observer surface itself. On the exact fixed-cutoff measurement surface it is generated by central record projectors, and practical readout may use approximately commuting projectors that are close to that central reference algebra. That is what makes the observer-facing record surface shareable across overlapping observer descriptions without violating the usual no-cloning constraints on generic quantum states.

12.2 Record algebras and definite outcomes

The integrated measurement appendices and Ref. [3] make the measurement interface precise at fixed cutoff. On the declared operational surface, the completed write/verify slice carries a finite commutative central record algebra. Practical readout may instead use projectors Q_a on the same declared slots with commuting central reference projectors $\widehat{Q}_a$, where

$$\delta_{\text{rec}} := \max_a \|Q_a - \widehat{Q}_a\|.$$

Then

$$\|[Q_a, Q_b]\| \le 4\,\delta_{\rm rec},$$

and, whenever $\|\widetilde{\rho} - \rho\|_1 \le \varepsilon$,

$$\left|\mathrm{Tr}(\widetilde{\rho} Q_a) - \mathrm{Tr}(\rho \widehat{Q}_a)\right| \le \varepsilon + \delta_{\rm rec}.$$

On the explicit fixed-cutoff screen architecture, the exact record algebra after a completed write/verify cycle is

$$\mathcal{Z}_{\rm rec}(t) = \mathrm{Alg}\Big(\{P_{m_I^{(J)}(t)}\}_{I\neq J}, \{P_{r_I^{\rm bulk}(t)}\}_I, \{\Pi_\alpha^{(IJ)}(t)\}_{I<J,\alpha}\Big).$$

Ref. [3] proves that, on the declared operational measurement surface, $\mathcal{Z}_{\rm rec}(t)$ is commutative and central for the readout instrument being used.

This fixed-cutoff centrality result is what supports definite outcomes without adding a second ontology. The integrated measurement ledger phrases the same idea through edge-center decomposition: on a fixed collar with exact Markov structure, the state splits into superselection blocks,

$$\rho_{ABC} = \bigoplus_j q_j\, \rho^{(j)}_{Ab_L^{(j)}} \otimes \rho^{(j)}_{b_R^{(j)}C},$$

and the label j is classical center data. The supplement then goes one step further: in the physical algebra there are no interference observables between different sector blocks. So once a particle detection event has been recorded in the observer-accessible center/record algebra, the accessible physics is organized as a classical mixture over those blocks rather than as a superposed measurement record.

Accordingly, a measured particle family or excitation channel should be thought of as an observer-accessible sector or record value carried by the central measurement surface, not as a mysterious absolute collapse of the full universe-state into a preferred basis chosen by hand.

12.3 Born probabilities and post-measurement states

The fixed-cutoff measurement package is also explicit about probabilities and updates. For every event E in the sigma algebra generated by $\mathcal{Z}_{\rm rec}(t)$, the measurement probability is

$$\mathbb{P}_t(E) = \mathrm{Tr}\big(\rho_t P_E\big),$$

where $P_E \in \mathcal{Z}_{\rm rec}(t)$ is the projector for that event. Conditioning on the event then produces the post-measurement state

$$\rho_t|_E = \frac{P_E \rho_t P_E}{\mathrm{Tr}(\rho_t P_E)}.$$

In the supplement, the same statement appears as the Lüders update on the record algebra. In the synthesis paper *Observers Are All You Need* [1], this fixed-cutoff Born/Lüders package and its Bell/CHSH extension are imported as theorem-bearing rather than non-theorem commentary.

This directly answers the particle-state question in the present framework. Measurements give particles actual states by conditioning the observer-accessible state on the central record algebra. A detector click, a stable overlap-sector readout, or a pointer value does more than announce a pre-existing classical label; it defines the conditioned state for subsequent observer-accessible physics. If the same event is re-read without any accepted repair move touching its support, the result in Ref. [3] shows that the re-read probability is 1. So the measurement interface is well defined and operationally auditable.

12.4 Particle identity as transport-stable structure

Once one asks "which event happened?" and then "which particle family is this?", the flavor derivation becomes essential. The active flavor derivation does not start with the names electron, muon, tau lepton, or up quark built into the fundamental observable. It starts with transport kernels, generation-bundle data, same-label eigenline transport, overlap-edge cocycles, and a persistent flavor observable carrying intrinsic labels $f1, f2, f3$. This is an important discipline condition in the paper. At the deepest flavor surface used here, family identity is transport-stable intrinsic structure first and named experimental family assignment second.

The shared Yukawa/excitation dictionary carries substantial weight even though it is not itself the leading missing object. The flavor pipeline exports an invariant base: projectors, spectral gaps, pair suppressions, cycle phases, and common sector-response objects. The nonpromoting component is the map from that intrinsic base to the named low-energy fermion families when the required sector bridge is not emitted. The quark derivation emits named rows because the forward branch is far enough along to support direct numerical comparison there; the charged-lepton derivation does not do so at theorem grade, and the neutrino derivation does so only on its weighted-cycle theorem branch rather than on a fully flavor-labeled closure surface, even though they use the same downstream family machinery.

Accordingly, particle identity is a refinement-stable observer-accessible transport pattern whose named interpretation is inherited from the available dictionary and comparison surfaces. Where that dictionary is work in progress, the paper states that boundary explicitly.

12.5 Scope boundary of the measurement chapter

This whole chapter has to preserve one final distinction. The fixed-cutoff measurement interface is theorem-bearing on the microphysics surface: central record algebra, Born probabilities, Lüders conditioning, repeated-read stability, and the stated two-wing Bell/CHSH theorem stack are written there as fixed-cutoff statements. What is *not* proved at the same level is the stronger global observer-continuation or strange-loop proposal sometimes discussed in broader OPH notes. This paper should therefore use the fixed-cutoff measurement package directly, while leaving the larger metaphysical closure proposal out of the main technical chain.

13 Discussion: Matter, Antimatter, Supersymmetry, and Continuation Questions

This section is intentionally not written in the same voice as the structural and quantitative sections above. The OPH theorem surface distinguishes sharply between recovered-core results, quantitative-closure branches, secondary quantitative branches, and continuations. Matter versus antimatter, supersymmetry, and several adjacent questions lie on that boundary. They are too important to omit from a particle-spectrum paper, but they are not part of the closed prediction surface in the same way that the bosonic and quark rows are.

13.1 Matter and antimatter

The Standard Model structural branch fixes the chiral gauge architecture within which matter and antimatter are defined. Once a chiral fermion representation is present, its conjugate representation gives the corresponding antiparticle content. In that limited sense, OPH explains why matter and antimatter both exist: they are paired by the realized gauge and chiral structure of the low-energy branch.

What is *not* structurally closed is the cosmological asymmetry between them. The integrated continuation notes record a baryogenesis continuation tied to the realized $\mathbb{Z}_6$ quotient and the electroweak topological channel. In that continuation, the entropy deficit of the $\mathbb{Z}_6$ quotient gives a suppression factor

$$\varepsilon = e^{-\ln 6} = \frac{1}{6},$$

and the same notes compare powers of that factor with the 12-fermion electroweak baryon-violating channel. But that comparison is only a suppression-counting heuristic. It is not a dynamical derivation of the observed baryon asymmetry.

But the observer-side status notes are explicit about the status of that estimate: the baryon asymmetry scale is a Phase-III continuation, and suppression-counting estimates are not a substitute for a derived out-of-equilibrium mechanism. The declared derivation also does not expose a baryogenesis output row. So this paper can report the available OPH continuation idea, but it should not present a sharp matter/antimatter abundance prediction as if it were a closed theorem or an emitted output.

The required ingredients are specific: a concrete out-of-equilibrium epoch, justified defect/sphaleron production dynamics, derived CP-odd source terms, transport and washout control, and a freeze-out computation of the final asymmetry. The available corpus does not have a baryogenesis theorem package. The same continuation also notes that CP-odd overlap data are generically available because central overlap phases are sent to their inverses under CP, but that observation is not a baryogenesis theorem.

13.2 Why the declared branch does not require supersymmetry

The D10 appendix below contains the dedicated gauge-coupling quantitative-closure discussion unification. Its main point is not that supersymmetry has been mathematically ruled out in every possible extension of the program. Its point is narrower and more relevant to this paper: the declared branch does not *need* a supersymmetric partner spectrum in order to explain why unification-like running might appear.

At one loop, the integrated D10 appendix records the familiar fact that Standard Model beta-function coefficients do not produce successful naive unification, whereas MSSM-like coefficients do. The OPH continuation idea is then that edge-sector heat-kernel weights can shift the effective beta-function coefficients in an MSSM-like direction through geometric or entropic sector multiplicity rather than through an actual low-energy superpartner spectrum. In that reading, "unification-like behavior" and "a supersymmetric particle zoo" come apart.

That continuation is narrower than a derived unification theorem. The appendix records one calibration branch in which Peter–Weyl multiplicities, the printed one-loop running frame, the threshold conventions, and an additional fermionic-grading restriction together produce an MSSM-like benchmark shift. It does not prove that those effective loop multiplicities, statistics restrictions, and decoupling conventions are uniquely forced by OPH edge sectors alone.

For this paper, the correct conclusion is therefore scope-limited and modest:

1. the derivation reconstructs the realized Standard Model branch without having to introduce a supersymmetric partner for every known field;
2. the integrated D10 appendix contains a continuation-level argument that some unification-style running features could arise from edge-sector structure rather than MSSM particle content;

3. the beta-shift package itself sits on explicit D10 calibration assumptions rather than a closed edge-sector theorem;
4. none of this is a universal theorem that supersymmetry is impossible.

So if the paper asks "why is there no supersymmetry in the derived spectrum?", the best technical answer is: because the realized branch stops at the Standard Model content, and the existing unification discussion is trying to reproduce the relevant running behavior through a separate D10 calibration package, without extending the particle content to a full supersymmetric multiplet structure.

13.3 Three generations, flavor closure, and work in progress

The structural Standard Model branch fixes $N_g = 3$ on the realized MAR-admissible branch. This result is separate from the harder flavor closure problem. Three generations are structurally recovered; the full flavor dictionary, excitation map, and CKM closure are work in progress, while the charged source landing from P to physical charged data is a corpus-limited no-go and the neutrino lane is closed only on its declared weighted-cycle theorem branch rather than on a fuller flavor-labeled closure surface. The flavor chapter and the later family chapters therefore spend substantial space on constructive object boundaries as well as final numbers.

13.4 Ordinary matter and the hadron backend boundary

A second distinction belongs in discussion form rather than in the results sections: the difference between "elementary rows exist" and "ordinary matter is fully derived." Ordinary visible matter is dominated by hadrons, especially protons and neutrons. The OPH derivation has a serious hadron pipeline, with source-only hadron masses gated by a production backend export. The route toward ordinary matter is architecturally present, and the full observed matter spectrum is work in progress.

13.5 How to read these continuation questions

The reader should therefore interpret this whole discussion section with the same rule used by *Recovering Relativity and Standard Model Structure from Observer Overlap Consistency* [4]. A failure in one of these continuation branches would not by itself undo the recovered-core gauge and gravity chain, nor would it erase the reported bosonic and quark rows. Charged-lepton source landing, fuller flavor-labeled neutrino, and source-only hadron mass theorems sit beyond the closed structural core. Their truth or failure does not alter the recovered-core gauge and gravity chain, and it does not erase the reported bosonic and quark outputs.

14 Conclusion

The particle-spectrum derivation is nontrivial and much more concrete than a qualitative aspiration. The structural theorem surface fixes the realized Standard Model branch and the exact massless carrier sectors. The completion pipeline closes or scopes four non-hadron chains: structural massless carriers, the Higgs/top row on the declared D10/D11 surface, selected-class quarks on f_P, and neutrino absolute attachment on the weighted-cycle branch. The ledger also marks four separate row classes: the candidate P-closure source audit, the electroweak W/Z validation row, the charged-lepton absolute-mass landing, and the hadron backend plus empirical closure split. The W/Z row is a frozen compare-only adapter with its own public validation records. The Ward-projected Thomson lane supplies the fine-structure endpoint on the declared P-closure surface. In the charged lane the theorem surface carries an exact same-family witness, a determinant-line lift, and an algebraic mass readout from theorem-grade $A_{\mathrm{ch}}(P)$, while the source landing from the D10 quantitative descendants of P to physical charged data is closed as a corpus-limited no-go. Source-only hadron masses require a working OPH hadron backend. Empirical hadron closure values use a separate $e^+e^- \to$ hadrons payload class.

That accounting belongs on the page because it is part of the derivation. The paper puts the structural results, the fixed-point fine-structure value, the emitted quantitative rows, and the status table onto one common surface. The bosonic P-driven trunk is explicit. The local extension surface identifies the shared edge entropy with the D10 nonabelian sum on the lifted product presentation of the realized quotient branch, and the same surface fixes that scalar to $P/4$. The particle spectrum is therefore organized by one local screen scale rather than by a list of unrelated particle inputs. The status table records scope boundaries for the weak validation pair, charged-lepton witnesses, direct-top auxiliary codomain, strong-CP phase status, neutrino comparison tension, source-only hadron backend rows, and empirical hadron closure rows.

A D10 Heat-Kernel Calibration Supplement

This appendix carries the D10 compact-group heat-kernel items needed by the particle surface. The theorem-bearing leaf for the same-overlap thermal/Casimir law sits on the screen-microphysics paper surface at fixed cutoff. What is recorded here is the compact-group / Peter–Weyl lift and beta-shift bookkeeping used by the particle-side calibration branch. The public D10 status is explicitly Phase II rather than recovered-core Phase I. In particular, this appendix does not prove that OPH derives the physical one-loop beta coefficients from edge sectors alone.

A.1 One-Loop Calibration Frame

At one loop, if couplings unify at (M_U, α_U),

$$\alpha_i^{-1}(M_Z) = \alpha_U^{-1} + \frac{b_i}{2\pi} \ln \frac{M_U}{M_Z}.$$

On the D10 branch these are the branch values $(M_U(P), \alpha_U(P))$ emitted by the forward D10 solve; they are not inferred by inverse readback from measured low-energy couplings. Writing $A_i := \alpha_i^{-1}(M_Z)$ and $L := \ln(M_U/M_Z)$, one finds

$$L = \frac{2\pi}{b_1 - b_2}(A_1 - A_2),$$

and the corresponding consistency relation for the D10 lane is

$$A_3^{\mathrm{pred}} = \frac{b_3 - b_2}{b_1 - b_2} A_1 + \frac{b_1 - b_3}{b_1 - b_2} A_2.$$

This appendix keeps that algebra on the page as a calibration relation, not as a theorem that OPH derives a supersymmetric UV spectrum.

A.2 Benchmark coefficient comparison

For the standard one-loop benchmark coefficients,

Model	(b_1, b_2, b_3)	$\alpha_s(M_Z)^{\text{pred}}$
SM	$(41/10, -19/6, -7)$	≈ 0.071
MSSM	$(33/5, 1, -3)$	≈ 0.116
Observed		0.1179 ± 0.0010

At one loop, the printed MSSM-style coefficients reproduce the observed closure far better than the plain SM coefficients. On this D10 branch this is a benchmark matched by the D10 calibration branch, not a theorem that OPH derives a supersymmetric UV spectrum.

A.3 Heat-Kernel Sector Law

Read this subsection as the compact-group merge boundary above the fixed-cutoff microphysics theorem. MaxEnt plus the bi-invariant compact-group package yield the familiar $d_R e^{-tC_2(R)}$ law once the Peter–Weyl lift is supplied; they are not a second, independent proof of the finite-cutoff same-overlap Casimir branch.

Theorem A.1 (Heat-kernel edge-sector weights)**.** *Under MaxEnt with bi-invariant constraints on a compact Lie group G, the sector probabilities take heat-kernel form:*

$$p_R(t) \propto d_R\, e^{-tC_2(R)},$$

where $d_R = \dim R$ and $C_2(R)$ is the quadratic Casimir.

Lemma A.2 (Bi-invariant operators)**.** *If $G = \prod_i G_i$ is a compact semisimple Lie group written as a product of compact simple factors, then any bi-invariant second-order differential operator on G has the form*

$$D = c_0 \mathbf{1} - \sum_i c_i \Delta_{G_i},$$

where Δ_{G_i} is the Laplace–Beltrami operator on the factor G_i.

Remark A.3. *Bi-invariance is equivalent to $D \in Z(U(\mathfrak{g}))$. Linear terms vanish because there are no invariant vectors in the adjoint representation, and the quadratic part is proportional to the Killing form on each simple factor, giving one independent coefficient per factor. The Casimir element then acts as $-\Delta_G$ in the regular representation.*

A.4 Peter–Weyl Multiplicity and Beta Shifts

By Peter–Weyl decomposition, the effective refinement-limit edge representation space is

$$L^2(G) \cong \bigoplus_R V_R \otimes V_R^*.$$

Entanglement traces over one factor give multiplicity d_R in p_R, but loops see both factors, so the effective multiplicity entering the calibration branch is

$$N_{\text{eff}}(R) = d_R\, p_R.$$

The one-loop beta shift from edge sectors is then

$$\Delta b_a = \sum_R p_R\, d_R\, T_a(R),$$

where $T_a(R)$ is the Dynkin index for gauge factor a. This is the exact mathematical point on the D10 lane where edge-sector heat-kernel weights feed the calibration branch. Matching to MSSM-like one-loop running is a declared Phase II benchmark, not a promoted theorem-level MSSM spectrum claim.

The branch boundary for theorem-level unification closure is explicit. One requires a derived refinement-limit identification of which transportable sectors actually contribute to the running carrier, a derived statistics/grading rule rather than an imposed fermionic-grading restriction, controlled threshold and decoupling conventions on that same carrier, a derived representation content/truncation rule for the sectors retained in the comparison, and a proof that the effective loop multiplicities entering Δb_a are exactly the ones used in this calibration package rather than nearby alternatives. On the declared runtime surface the RG/matching/threshold/scheme packet is work in progress: scheme lock, threshold map, beta provenance, and interval composition.

At the branch unification diffusion parameter $t_U(P) \approx 1.64$ selected by the forward D10 solve, the benchmark shift is

$$\Delta b \approx (2.49, 4.38, 3.97) \qquad \text{vs MSSM} \qquad (2.50, 4.17, 4.00).$$

With the additional D10 calibration assumption of a fermionic-grading restriction to half-integer SU(2) sectors, this sharpens to

$$\Delta b \approx (2.50, 4.17, 3.97),$$

which matches the MSSM-style one-loop benchmark at the percent level under that declared calibration package. Equivalently, the benchmark ratio

$$\frac{\Delta b_3}{\Delta b_2} \approx 0.91$$

sits close to the MSSM comparison value 0.96.

References

[1] B. Müller, A. Osika, K. Xue, B. Cassie, P. Nguyen, M. Poneder, and K. A. Anirudha, *Observers Are All You Need*. GitHub PDF: https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/observers_are_all_you_need.pdf

[2] B. Müller, K. Xue, and K. A. Anirudha, *Reality as a Consensus Protocol: The Fixed-Point Computation That Implements Physics*. GitHub PDF: https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/reality_as_consensus_protocol.pdf

[3] B. Müller, A. Osika, K. Xue, and B. Cassie, *Federated Echosahedral Screen Microphysics: Patch Hardware, Records, and Observer Synchronization in OPH*. GitHub PDF: https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/screen_microphysics_and_observer_synchronization.pdf

[4] B. Müller, A. Osika, K. Xue, and P. Nguyen, *Recovering Relativity and Standard Model Structure from Observer Overlap Consistency*. GitHub PDF: https://github.com/FloatingPragma/observer-patch-holography/blob/main/paper/recovering_relativity_and_standard_model_structure_from_observer_overlap_consistency_compact.pdf

[5] National Institute of Standards and Technology, *2022 CODATA Recommended Values of the Fundamental Constants of Physics and Chemistry*, NIST SP 959, May 2024. https://physics.nist.gov/cuu/pdf/wallet_2022.pdf

[6] Planck Collaboration, *Planck 2018 results. VI. Cosmological parameters*, Astronomy & Astrophysics 641, A6, 2020. https://arxiv.org/abs/1807.06209

[7] DESI Collaboration, *Constraints on Neutrino Physics from DESI DR2 BAO and DR1 Full Shape*, arXiv:2503.14744, 2025. https://arxiv.org/abs/2503.14744

Sector	Claim tier	Derivation stage	Public output(s) / exact sidecar(s)	Caveat / theorem boundary
Structural carriers	structural	gauge/gravity structural theorems	exact $m_\gamma = m_g = m_{\rm grav} = 0$	theorem-grade structural exactness; no particle-side caveat
Electroweak bosons	compare-only adapter	D10 electroweak branch plus frozen validation adapter	$M_W = 80.377$ GeV, $M_Z = 91.18797809193725$ GeV	compare-only validation beneath the declared D10/P/fine-structure arithmetic bridge; public validation records are source spectral measure payload, same-scheme remainder, and interval certificate
Higgs/top stage	exact source-only split theorem on the declared D10/D11 running, matching, and threshold surface + compare-only exact sidecar	D10 gauge core followed by the target-free D10 repair chart, the shared scalar $\rho_{HT} = \log(1 + \tau_{2,\rm tree}^{\rm exact})$, the source-only split selectors, and the D11 Jacobian readout	exact pair $m_H = 125.1995304097179$ GeV, $m_t^{D11} = 172.3523553288312$ GeV; exact inverse-sidecar pair on the D11 Jacobian	Higgs row is closed on the declared surface. The inverse pair is compare-only. The auxiliary direct-top PDG row is a distinct compare-only codomain with a corpus-limited no-go boundary.
Quark family	selected-class theorem + exact supporting surfaces	shared excitation dictionary → D12 mass bridge → selected public quark frame class f_P → direct public sigma-datum descent → affine mean law → ordered three-point readout → exact forward construction	exact running quark sextet on the selected public quark frame class f_P, together with explicit exact forward Yukawas Y_u, Y_d on that selected class	selected-class closure only. The exact sextet matches the official PDG 2025 API running-quark surface, with the top coordinate taken from the PDG cross-section entry. The auxiliary direct-top PDG row is compare-only with a corpus-limited no-go boundary. The stronger class-uniform/global public-frame classification is not emitted; this is a corpus-limited no-go boundary without weakening the selected-class theorem. Strong CP is work in progress: the available corpus does not derive $\theta_{\rm QCD}$, does not emit physical $\bar\theta$, and does not prove that the physical strong-CP phase vanishes. Supporting exact surfaces: same-family witness on `current_family_only`; restricted common-refinement transport-frame theorem on `current_family_common_refinement_transport_frame_only`. Both realize the same sextet and explicit exact forward Yukawas. Separate target-free mass bridge: $\Delta_{ud}^{\rm overlap} = \frac{1}{6}\log(c_d/c_u)$, equivalently $\Theta_{ud}^{\rm mass}$ = `quark_same_family_value_law`, on the emitted D12 mass ray. Lower selector: $\sigma_{\rm ref}$ is a negative same-label sheet selector statement beneath the selected-class theorem. No global classification of all quark frame classes is claimed
Charged leptons	continuation + exact sidecar witness	shared excitation dictionary → ordered charged carrier → exact centered readback → determinant-line lift on theorem-grade physical charged data	exact same-carrier centered readback together with the exact same-family charged triple on a target-anchored witness	public charged masses are not emitted from P. The theorem lane does not emit a theorem-grade sector-isolated charged determinant exponent vector. It does not attach a source-side determinant character to the physical charged determinant line. The determinant-line lift and algebraic mass readout apply only on theorem-grade physical charged data
Neutrinos	continuation + emitted bridge invariant + diagnostic sidecars	same-label scalar certificate → weighted-cycle branch → bridge rigidity → absolute attachment → shared-basis transport readout	emitted absolute masses, central splittings, and the physical Majorana pair $(\alpha_{21}^{(\rm Maj)}, \alpha_{31}^{(\rm Maj)}) = (153.618518^\circ, 257.003241^\circ)$ on the weighted-cycle branch, with bridge invariant C_ν above emitted proxy P_ν and paper-facing amplitude B_ν	theorem-grade on the declared weighted-cycle branch; the exact positive-segment adapter, bridge corridor, and correction audit are diagnostic-only and do not feed back into theorem state
Hadrons	source backend absent; empirical closure policy emitted	Ward-projected spectral-measure contract for an OPH hadron backend, plus empirical $e^+e^- \to$ hadrons source registry and payload schema	no source-only hadron masses; empirical hadron closure rows are separate from source-only OPH rows	Source-only hadron prediction requires a working OPH hadron backend, such as GLORB/Echosahedron, with the Ward-projected hadronic spectral measure and systematics. The empirical closure surface uses a separate $e^+e^- \to$ hadrons payload class and cannot promote the source-only theorem

Lane	Exact output(s)	Exact chain on the paper surface	Caveat
Structural carriers	$m_\gamma = m_g = m_{\text{grav}} = 0$	axioms → realized electromagnetic/color/dynamical-metric carrier skeleton → symmetry-protected zeros	theorem-grade structural exactness
Electroweak exact sidecar	exact frozen W/Z pair	axioms → $D7$ → $D10$ quantitative-closure chain → exact frozen repair pair `oph_d10_ew_w_anchor_neutral_shear_factorization`	exact only on the frozen authoritative repair surface; compare-only beneath the target-free theorem
Higgs/top exact sidecar	exact Higgs/top inverse pair	axioms → $D10$ gauge core → $D11$ Jacobian → exact inverse slice	compare-only inverse slice; the public Higgs/top pair comes from the source-only split theorem on the declared surface, while the repo-wide exact public top row is also carried by the selected-class quark theorem
Charged exact witness	exact same-family charged triple	axioms → shared excitation dictionary → ordered charged carrier → exact centered readback → closed quadratic readout theorem → same-family exact witness	same-family-only; theorem lane carries exact centered readback plus the closed common-shift no-go. The available corpus has a no-go boundary: a theorem-grade $\widehat{C}_e$ promotion and uncentered trace lift emitting the physical scalar $\mu_{\text{phys}}(Y_e)$ are not part of the emitted theorem surface
Quark selected-class theorem	exact running quark sextet plus explicit exact forward Yukawas on f_P	axioms → shared excitation dictionary → D12 mass bridge → selected public quark frame class f_P → direct public sigma-datum descent → affine mean law → ordered three-point readout → exact forward construction	selected-class closure on the public physical quark frame class chosen by P; the sextet matches the official PDG 2025 API running-quark surface exactly, with the top coordinate taken from the PDG cross-section entry; the auxiliary direct-top PDG row is compare-only with a corpus-limited no-go boundary; the stronger global classification of public quark frame classes is not emitted and has a corpus-limited no-go boundary; the same exact sextet is also realized by the same-family witness on `current_family_only` and by the restricted common-refinement transport-frame exact chain, which also carries explicit exact forward Yukawas Y_u and Y_d; the target-free mass bridge $\Delta_{ud}^{\text{overlap}} = \frac{1}{6}\log(c_d/c_u)$, equivalently $\Theta_{ud}^{\text{mass}} =$ `quark_same_family_value_law`, closes separately on the emitted D12 mass ray; the lower selector σ_{ref} is a negative sheet statement beneath the selected-class theorem; no global classification of all quark frame classes is claimed
Neutrino theorem branch	emitted absolute mass triple, representative central splittings, and physical Majorana pair	axioms → same-label scalar certificate → weighted-cycle branch → bridge-rigidity theorem → absolute-attachment theorem → shared-basis transport readout	theorem-grade on the declared weighted-cycle branch; the exact positive-segment adapter and bridge-coordinate sidecars are diagnostic-only

Carrier	Claim tier	OPH output	Physical role	Why the mass is zero
Photon	structural	0 GeV	electromagnetic force carrier	gauge redundancy forbids a hard mass
Gluons	structural	0 GeV	color-force carriers	color-gauge redundancy forbids a hard mass
Graviton	structural	0 GeV	spin-2 metric carrier	diffeomorphism redundancy forbids a hard mass

Particle	Stage	Claim tier	Report
W boson	D10 electroweak branch	compare-only frozen adapter	80.377
Z boson	D10 electroweak branch	compare-only frozen adapter	91.1879
Higgs boson	Higgs/top critical stage	quantitative theorem	125.199
Top quark	Higgs/top critical stage	quantitative theorem on declared surface; direct-top auxiliary codomain no-go	172.352

Quark	Claim tier	Exact running value	Derivation stage
Up quark	selected-class theorem	0.00216 GeV	selected public quark frame class f_P
Down quark	selected-class theorem	0.00470 GeV	selected public quark frame class f_P
Strange quark	selected-class theorem	0.0935 GeV	selected public quark frame class f_P
Charm quark	selected-class theorem	1.273 GeV	selected public quark frame class f_P
Bottom quark	selected-class theorem	4.183 GeV	selected public quark frame class f_P
Top quark	selected-class theorem	172.3523553288311 GeV	selected public quark frame class f_P