Structural Closure of the Coupled Dirac–Λ Framework on the Adopted S^3 Carrier: Global Mass Determination and Scheme Rigidity

Jeremy Rodgers

April 24, 2026

Abstract

We prove two closure theorems for the coupled Dirac–Λ framework on a fixed compact Riemannian carrier. For the mass determination problem, we establish: (i) the spectral action is strictly concave in the Yukawa mass variables on $\mathbb{R}^9_{>0}$, with diagonal negative-definite Hessian; (ii) for any carrier M satisfying the dominance margin condition $\Delta(M) > 1$, the physical mass configuration is the unique global maximum of the spectral action subject to the capacity inequality, giving unrestricted global uniqueness of the fermion mass vector on that carrier; (iii) the three anchor masses used in the companion paper are self-consistently determined by the overdetermined KKT system (20 equations for 18 unknowns), eliminating all observational mass input; (iv) all four active scales are intrinsically determined by structural conditions (sensitivity extrema, DN-threshold crossing, and saturation peak condition), removing the sole remaining tolerance parameter. For the scheme rigidity problem, we establish: (v) the spectral filter is constrained to the family $g_p(u) = (1 + u^2)^{-p/2}$ with $p \geq 6$ by shell-sum convergence, moment sign conditions, and capacity-gradient sign structure; (vi) the boundary budget D_T is unique up to a bounded correction (collecting the result from companion papers); (vii) the effective band is functionally determined by the filter; (viii) all structural predictions (generation number, gauge trace ratios, hierarchy, cascade, $\theta = 0$) are invariant under admissible scheme variation, while exact mass values have bounded smooth dependence on the scheme.

These results establish that on any compact Riemannian carrier M with $\Delta(M) > 1$, the nine SM fermion masses are uniquely determined on that carrier by the carrier geometry, with no free continuous parameters remaining in the Tier-1 mass sector. The mass values are carrier-specific: exact Yukawa ratios depend on the geometric deficit $G_M(T)$ and are not universal law-level outputs. This geometry-dependence is a structural boundary of the framework, established by independent no-go theorems both within E1–E8 [14] and at the Tier-0 lawhood level [15]. The structural properties—concavity, global uniqueness, hierarchy forcing, cascade ordering, scheme invariance—are geometry-independent and hold for all admissible carriers.

Contents

1 Introduction

The coupled Dirac–Λ framework, developed in the companion papers [1, 2, 3], establishes a scale-by-scale capacity inequality $S_\Omega(\mathcal{D}_E;T) \leq D_T(K_{\mathrm{OS}})$ between the spectral content of the bulk Dirac operator and the dissipative budget of the boundary operator K_{OS}. Within this framework, three generations are forced [1], the boundary operator is uniquely identified as the Dirichlet-to-Neumann map of the Standard Model spectral triple [2], and the fermion mass prediction problem is reduced to a finite-dimensional balance equation with verified structural properties [3].

Independent no-go theorems establish that exact Yukawa ratios are not universal law-level outputs: within E1–E8, the physical Yukawa quotient manifold retains nonempty interior in the viable set [14]; and at the Tier-0 lawhood level, the obstruction persists in any physically lawful framework [15]. The mass values are carrier-specific, depending on the geometric deficit $G_M(T)$. Within this context, two closure problems remained:

P1. Mass determination on a fixed carrier. Given a specific carrier M, are the nine fermion masses uniquely determined by the KKT saturation equations, and is the solution globally unique?

P2. Scheme rigidity. Is the globally fixed scheme data (filter, determinant prescription, effective band) uniquely determined?

The present paper resolves both problems. For P1, we prove global uniqueness of the mass vector on any carrier satisfying $\Delta(M) > 1$ by identifying the physical configuration as the unique maximum of a strictly concave function subject to the capacity constraint, and we eliminate all anchor masses by exhibiting an overdetermined self-consistent system. For P2, we classify all structural predictions into scheme-invariant (the majority) and scheme-dependent with bounded variation (the exact mass values), and we characterize the admissible filter family. Within the broader Everything Equation programme [13], these results sit at Layer 4 (interface realization), where the upstream Tier-0 Λ-selection chain has now resolved the β-closure problem and the framework derives $\alpha^{-1} = 137$ with zero free parameters on the adopted S^3 carrier.

Section 2 establishes the strict concavity of the spectral action in mass variables and the boundedness of the feasible set. Section 3 derives all active scales intrinsically and obtains the derived quadratic-moment bound. Section 4 proves unrestricted global uniqueness via cascade separation, LICQ, and the dominance inequality. Section 5 formulates the dominance margin as an intrinsic manifold filter. Section 6 eliminates the anchor masses. Section 7 constrains the spectral filter. Section 8 classifies scheme-invariant predictions. Section 9 states the combined closure theorem.

2 Strict Concavity of the Spectral Action in Mass Variables

2.1 The spectral action restricted to Yukawa variables

We work on the product geometry $M \times F$ with the Chamseddine–Connes finite spectral triple. Following [1, 3], the Yukawa-squared variables are $X_\alpha = m_\alpha^2/x^*$ for $\alpha = 1, \ldots, 9$, where $x^* = v^2/2$ and m_α are the nine fermion masses.

Definition 2.1 (Spectral action restricted to mass variables). *Fix the geometric background (M, g_E, A) and the scheme data $(f, g, \det_*, [T_{\mathrm{UV}}, T_{\mathrm{IR}}])$. The internal spectral action restricted to the Yukawa variables is*

$$\mathcal{S}^{(\mathrm{int})}(\mathbf{X}) = \sum_{\alpha=1}^{9} \mathrm{mult}(\alpha)\, P_{\mathrm{SA}}(X_\alpha), \tag{1}$$

where $P_{\mathrm{SA}}(X) = A_1 + 3A_3 x^{*2} X^2 + 4A_4 x^{*3} X^3$ is the spectral-action stationarity polynomial [1], with $A_1 > 0$, $A_3 < 0$, $A_4 < 0$. The geometric contributions to $\mathcal{S}$ are independent of $\mathbf{X}$ and are absorbed into the geometric deficit.

Theorem 2.2 (Strict concavity)**.** *The internal spectral action $\mathcal{S}^{(\mathrm{int})} : \mathbb{R}^9_{>0} \to \mathbb{R}$ is strictly concave.*

Proof. The Hessian of $\mathcal{S}^{(\mathrm{int})}$ is diagonal:

$$\frac{\partial^2 \mathcal{S}^{(\mathrm{int})}}{\partial X_\alpha \partial X_\beta} = \delta_{\alpha\beta}\, \mathrm{mult}(\alpha)\, P''_{\mathrm{SA}}(X_\alpha). \tag{2}$$

From $P_{\mathrm{SA}}(X) = A_1 + 3A_3 x^{*2} X^2 + 4A_4 x^{*3} X^3$:

$$P'_{\mathrm{SA}}(X) = 6A_3 x^{*2} X + 12 A_4 x^{*3} X^2, \tag{3}$$

$$P''_{\mathrm{SA}}(X) = 6A_3 x^{*2} + 24 A_4 x^{*3} X. \tag{4}$$

Since $A_3 < 0$ and $A_4 < 0$, both terms are strictly negative for all $X > 0$. Hence $P''_{\mathrm{SA}}(X) < 0$ on $\mathbb{R}_{>0}$, every diagonal entry of the Hessian (2) is strictly negative, and $\mathcal{S}^{(\mathrm{int})}$ is strictly concave on $\mathbb{R}^9_{>0}$. □

Remark 2.3. *Strict concavity implies uniqueness of any unconstrained maximizer. For the constrained problem (maximizing $\mathcal{S}^{(\mathrm{int})}$ subject to the capacity inequality), constrained uniqueness requires either a convex feasible set or injectivity of the KKT map; we establish the latter in section 4.*

Remark 2.4. *The strict concavity is a consequence of the capacity structure: the saturation kernel $h(u) = g(u)\,q(u)$ satisfying $h(0) = 0$ eliminates the A_2 moment, producing the depressed cubic P_{SA} with $A_3, A_4 < 0$. A generic Chamseddine–Connes cutoff with $A_2 \neq 0$ yields a full cubic that can be convex in some regions of $\mathbb{R}_{>0}$.*

2.2 Boundedness of the feasible set

Lemma 2.5 (Bounded feasible domain)**.** *The feasible set*

$$\mathcal{F} := \left\{ \mathbf{X} \in \mathbb{R}^9_{>0} : S_\Omega(\mathcal{D}_E(\mathbf{X}); T) \leq D_T(K_{\mathrm{OS}}(\mathbf{X}))\ \ \forall T \in [T_{\mathrm{UV}}, T_{\mathrm{IR}}] \right\} \tag{5}$$

is bounded.

Proof. In the UV regime ($T \to 0$), the capacity inequality gives $S^{(\mathrm{int})}_\Omega(T) \leq D^{(\mathrm{int})}_T(T) + G(T)$. The UV asymptotics $S^{(\mathrm{int})}_\Omega \sim (T/2) \sum_\alpha \mathrm{mult}(\alpha)\, m_\alpha^2$ and $D^{(\mathrm{int})}_T \sim (T/2) \sum_\alpha \mathrm{mult}(\alpha)\, m_\alpha$ give

$$\sum_\alpha \mathrm{mult}(\alpha)\, m_\alpha^2 \leq \sum_\alpha \mathrm{mult}(\alpha)\, m_\alpha + C_{\mathrm{UV}},$$

which bounds $\|\mathbf{m}\|^2$ and hence $\|\mathbf{X}\|$. □

Lemma 2.6 (Interior attainment)**.** *The maximum of $\mathcal{S}^{(\mathrm{int})}$ on $\overline{\mathcal{F}}$ is attained in the interior $\mathbb{R}^9_{>0} \cap \mathcal{F}$, not on the boundary $\{X_\alpha = 0\}$.*

Proof. Existence of an interior KKT point is established constructively in [3], where an explicit solution on S^3 is exhibited and shown to have full-rank Jacobian. Strict concavity of $\mathcal{S}^{(\mathrm{int})}$ then rules out boundary maximizers: at any boundary point $X_\alpha = 0$, the directional derivative $\partial \mathcal{S}^{(\mathrm{int})}/\partial X_\alpha|_{X_\alpha = 0} = 0$ (since $P'_{\mathrm{SA}}(0) = 0$), while the interior KKT point achieves a strictly higher value by virtue of satisfying the first-order stationarity conditions at a point where $P''_{\mathrm{SA}} < 0$. □

3 Intrinsic Scale Selection

The companion paper [3] derived four active scales, with the fourth (T_4) depending on a tolerance parameter $\varepsilon_{\mathrm{rec}}$. We now show all four scales are intrinsically determined.

3.1 Scales T_1 and T_2: sensitivity extrema

Definition 3.1 (Logarithmic sensitivity). *The logarithmic sensitivity of the capacity gradient is*

$$\sigma(T) := \frac{\partial}{\partial \log T} \operatorname{cap\_grad}(X; T, w_s)\Big|_{X=X_b}, \tag{6}$$

where X_b is a representative heavy-channel probe (e.g., the b-quark) [3].

Proposition 3.2 (Intrinsic T_1, T_2). *On S^3, the sensitivity $\sigma(T)$ has exactly two zeros in the physically relevant range, at $T_1/T^* \approx 0.0377$ (UV, Dirac-dominated) and $T_2/T^* \approx 5.01$ (mid-range, near DN crossover) [3]. These are intrinsic functionals of the spectral geometry of S^3 and the filter g. No external parameter is involved.*

3.2 Scale T_3: DN-threshold crossing

Proposition 3.3 (Intrinsic T_3). *The DN-threshold condition $\Theta_s(T_3) = T_3\sqrt{x^* X_s(T_3)} = 1$, where $X_s(T)$ is the tracked strange-quark branch, determines T_3 uniquely by the implicit function theorem [3]. This is the scale where the DN $k = 0$ argument exits the Taylor regime. No external parameter is involved.*

3.3 Scale T_4: saturation peak condition

Definition 3.4 (Saturation peak selector). *Define the saturation peak condition:*

$$T_4 \cdot x^* \cdot X_{\mathrm{top}}(T_4) = u^*, \tag{7}$$

where $u^* \approx 0.437$ is the unique maximum of the saturation kernel $h(u) = g(u)\, q(u)$ [3], and $X_{\mathrm{top}}(T)$ is the heavy root of the up-quark stationarity equation at scale T.

Theorem 3.5 (Intrinsic T_4). *The saturation peak condition* (7) *determines a unique scale T_4.*

Lemma 3.6 (Monotone plateau regime). *Define $\Psi(T) := T \cdot x^* \cdot X_{\mathrm{top}}(T) - u^*$. There exists $T_{\min} > 0$, determined by the filter tail bound $g_p(T_{\min} \cdot x^* \cdot X_{\mathrm{top}}^{(\infty)}) > 1/N_{\mathrm{diss}}$, such that Ψ is strictly increasing on $(T_{\min}, \infty)$.*

Proof. $\Psi'(T) = x^*[X_{\mathrm{top}}(T) + T\, X'_{\mathrm{top}}(T)]$. By implicit differentiation of $F_u(X; T, \mu(T)) = 0$ along the heavy root:

$$X'_{\mathrm{top}}(T) = -\frac{\partial_T F_u}{\partial_X F_u}.$$

For $T > T_{\min}$, the heavy root $X_{\mathrm{top}}(T)$ lies in the plateau regime where $X_{\mathrm{top}}(T) \to X_{\mathrm{top}}^{(\infty)} = m_t^2/x^* \approx 0.987$ and $\partial_T F_u$ vanishes as the capacity gradient saturates. Explicitly, $|T\, X'_{\mathrm{top}}(T)| \le C\, g_p(T \cdot x^* \cdot X_{\mathrm{top}}^{(\infty)})$ for a computable constant $C > 0$ depending on A_3, A_4, and w_q. The filter tail bound $g_p(u) = (1 + u^2)^{-3}$ ensures $|T\, X'_{\mathrm{top}}| < X_{\mathrm{top}}/2$ for $T > T_{\min}$, giving $\Psi'(T) > x^* X_{\mathrm{top}}/2 > 0$. □

Proof of theorem 3.5. For small T: $T \cdot x^* \cdot X_{\text{top}} < u^*$, so $\Psi(T) < 0$. For large T: $X_{\text{top}}(T)$ is bounded near m_t^2/x^* while $T \to \infty$, so $\Psi(T) \to +\infty$. By the intermediate value theorem, Ψ has at least one zero. By theorem 3.6, Ψ is strictly increasing on $(T_{\min}, \infty)$, so the zero in this regime is unique. Below $T_{\min}$, $\Psi < 0$ (since $T_{\min} \cdot x^* \cdot X_{\text{top}}(T_{\min}) < u^*$ by construction of $T_{\min}$). Hence the zero is globally unique. □

Corollary 3.7 (No tolerance parameter). *The four active scales $\{T_1, T_2, T_3, T_4\}$ are fully determined by intrinsic structural conditions. The tolerance parameter ε_{rec} from [3] is eliminated.*

Lemma 3.8 (Heavy-channel lower bound from intrinsic T_4). *At the intrinsic top scale T_4 satisfying the saturation peak condition* (7),

$$m_{\text{top}}^2 = x^* \, X_{\text{top}}(T_4) = \frac{u^*}{T_4} \;\geq\; \frac{u^*}{T_{\text{IR}}}, \tag{8}$$

since $T_4 \leq T_{\text{IR}}$ (by definition, active scales lie within the admissible window $[T_{\text{UV}}, T_{\text{IR}}]$; the saturation-peak solution is constructed within this window).

Proof. The saturation peak condition gives $T_4 \cdot x^* \cdot X_{\text{top}}(T_4) = u^*$, hence $x^* \, X_{\text{top}}(T_4) = u^*/T_4$. Since T_4 is an active scale lying within the admissible window $[T_{\text{UV}}, T_{\text{IR}}]$ by construction, $T_4 \leq T_{\text{IR}}$, giving the lower bound. □

Theorem 3.9 (Derived quadratic-moment lower bound). *Let* $\text{mult}_{\max} = 12$ *(the up-sector multiplicity). Then*

$$S_2 \;\geq\; \text{mult}_{\max} \, m_{\text{top}}^2 \;\geq\; 12 \, \frac{u^*}{T_{\text{IR}}}. \tag{9}$$

Using the electroweak crossing relation $v^2 = (2\pi)^2/T^*$ [3], this gives:

$$S_2 \;\geq\; c_2 \, v^2, \qquad c_2 := \frac{12 \, u^* \, T^*}{(2\pi)^2 \, T_{\text{IR}}}. \tag{10}$$

The constant c_2 is fully scheme-determined: $u^* \approx 0.437$ comes from the saturation kernel $h = g \cdot q$, T^* from the capacity crossing, and T_{IR} from the active scale window endpoint. No Standard Model mass values are used.

Proof. The first inequality follows from $S_2 = \sum_\alpha \text{mult}(\alpha) \, m_\alpha^2 \geq \text{mult}_{\max} \, m_{\max}^2 \geq 12 \, m_{\text{top}}^2$. The second uses theorem 3.8. For the conversion to v^2: the electroweak crossing relation identifies $E^* = 1/\sqrt{T^*} = v/(2\pi)$, hence $v^2 = (2\pi)^2/T^*$. Dividing (9) by v^2 gives (10). □

4 Unrestricted Global Uniqueness

4.1 The constrained maximization interpretation

Proposition 4.1 (Physical configuration as constrained maximum). *The physical mass configuration maximizes the spectral action subject to the capacity budget. The KKT system [1] consists of:*

(i) Primal feasibility: $\Phi(T; \mathbf{X}) \leq 0$ for all $T \in [T_{\text{UV}}, T_{\text{IR}}]$.

(ii) Dual feasibility: $\mu \geq 0$.

(iii) Complementary slackness: $\mu(\{T\}) > 0 \implies \Phi(T; \mathbf{X}) = 0$.

(iv) Stationarity: $\nabla_{\mathbf{X}}\mathcal{S} + \sum_k \mu_k \nabla_{\mathbf{X}}\Phi(T_k; \mathbf{X}) = 0$.

This is the necessary condition for a constrained maximum of the concave function $\mathcal{S}^{(\mathrm{int})}$ *subject to the semi-infinite inequality constraint* $\Phi \leq 0$.

Proof. The capacity inequality constrains the spectral content: $S_\Omega \leq D_T$. The spectral action $\mathcal{S}$ measures the total spectral content. The system maximizes spectral content within the dissipative budget, consistent with the hierarchy-forcing mechanism [3]: heavy masses are budget-costly, forcing most masses to be light. □

4.2 Cascade ordering and scale separation from hierarchy forcing

We derive cascade ordering and a quantitative heavy/light separation *without inserting Standard Model numerical masses*. The only inputs are: (i) the hierarchy-forcing mechanism [3], (ii) the intrinsic T_4 saturation peak selector (theorem 3.5), and (iii) the electroweak crossing relation [3].

The constant c_2 appearing below was derived in theorem 3.9 (§3) from the intrinsic T_4 selector. The cascade separation results of this subsection are used only downstream in §4.4 and §4.5.

Definition 4.2 (Weighted quadratic Yukawa moment)**.** *Define the weighted quadratic moment*

$$S_2 := \sum_{\alpha=1}^{9} \mathrm{mult}(\alpha)\, m_\alpha^2, \tag{11}$$

where the multiplicities $\mathrm{mult}(\alpha) \in \{4, 12\}$ *are fixed by the internal spectral triple and* $N_{\mathrm{diss}} = \sum_\alpha \mathrm{mult}(\alpha) = 84$.

Remark 4.3. *The lower bound* $S_2 \geq c_2\, v^2$ *with explicit* $c_2 > 0$ *is derived in theorem 3.9 from the intrinsic* T_4 *selector and the electroweak crossing relation. No assumption on the Yukawa/Higgs sector is needed; the bound follows from the framework's own scale selectors.*

Lemma 4.4 (Existence of a heavy channel)**.** *The maximal charged fermion mass*

$$m_{\max} := \max_{\alpha=1,\ldots,9} m_\alpha$$

satisfies the structural lower bound

$$m_{\max} \;\geq\; \sqrt{\frac{c_2}{N_{\mathrm{diss}}}}\, v, \tag{12}$$

where $c_2 = 12\, u^* T^*/((2\pi)^2\, T_{\mathrm{IR}})$ *is the derived constant from theorem 3.9.*

Proof. Since $m_\alpha^2 \leq m_{\max}^2$ for all α,

$$S_2 = \sum_\alpha \mathrm{mult}(\alpha) m_\alpha^2 \leq \left(\sum_\alpha \mathrm{mult}(\alpha)\right) m_{\max}^2 = N_{\mathrm{diss}}\, m_{\max}^2.$$

Combine with $S_2 \geq c_2 v^2$ (theorem 3.9) to obtain (12). □

Theorem 4.5 (Cascade separation without SM data)**.** *Assume:*

(i) Hierarchy forcing [3, Theorem 13.1] produces a critical mass m_{crit}, defined as the threshold below which the capacity gradient changes sign. Explicitly, $m_{\mathrm{crit}}^2 = u_{\mathrm{sign}}/T^*$ where u_{sign} is the unique zero of $h'(u)$ on $\mathbb{R}_{>0}$ (scheme-determined). At any feasible stationary point, at least $N_{\mathrm{light}} \geq 6$ of the nine charged masses satisfy $m_\alpha \leq m_{\mathrm{crit}}$.

(ii) The derived quadratic-moment bound $S_2 \geq c_2\, v^2$ (theorem 3.9) holds.

Then there exists a heavy/light separation ratio

$$\frac{m_{\max}}{m_{\mathrm{crit}}} \;\geq\; R_0 := \sqrt{\frac{c_2}{N_{\mathrm{diss}}}}\,\frac{v}{m_{\mathrm{crit}}}. \tag{13}$$

In particular, whenever $R_0 > 2$, the matched scales

$$T_\alpha^* := \frac{u^*}{m_\alpha^2} \tag{14}$$

are strictly ordered by mass:

$$m_\alpha < m_\beta \quad\Longrightarrow\quad T_\alpha^* > T_\beta^*,$$

and the scale set $\{T_\alpha^*\}$ exhibits nontrivial separation between at least one heavy channel and the light channels.

Proof. The ratio bound (13) follows from theorem 4.4. The scale ordering follows immediately from (14). □

Corollary 4.6 (Diagonal dominance from structural separation)**.** *Assume the hypotheses of theorem 4.5 and $p \geq 6$. If $R_0 > 2$, then at any active scale T_k matched to a light mass $m_k \leq m_{\mathrm{crit}}$, every heavier mass $m_\alpha \geq m_{\max} \geq R_0 m_{\mathrm{crit}} \geq R_0 m_k$ satisfies*

$$T_k m_\alpha^2 \;=\; u^* \left(\frac{m_\alpha}{m_k}\right)^2 \;\geq\; u^* R_0^2,$$

and hence the Dirac-side contributions to the active constraint Jacobian row are suppressed by

$$g_p(T_k m_\alpha^2) \;\lesssim\; (u^* R_0^2)^{-p}.$$

Consequently the off-peak suppression inequality (polynomial order $2p$) implies strict row-wise diagonal dominance of the active constraint Jacobian after cascade ordering, and therefore LICQ.

Remark 4.7. *The condition $R_0 > 2$ is a purely structural separation requirement. It does not invoke PDG masses. It is verified whenever the derived constant $c_2 = 12\, u^* T^*/((2\pi)^2\, T_{\mathrm{IR}})$ (theorem 3.9) satisfies $c_2 > 4\, N_{\mathrm{diss}}\, (m_{\mathrm{crit}}/v)^2$, which holds for any hierarchy-forcing critical mass m_{crit} well below v.*

4.3 The LICQ condition

Definition 4.8 (Active constraint Jacobian)**.** *At the KKT solution* $\mathbf{X}^*$ *with active set* $\mathcal{A} = \{T_1, \ldots, T_{|\mathcal{A}|}\}$*, define the active constraint Jacobian* $J \in \mathbb{R}^{|\mathcal{A}| \times 9}$ *by*

$$J_{k\alpha} := \frac{\partial \Phi}{\partial X_\alpha}(T_k; \mathbf{X}^*) = \mathrm{mult}(\alpha)\, \mathrm{cap\_grad}(X^*_\alpha; T_k, w_{s(\alpha)}). \tag{15}$$

Lemma 4.9 (Off-peak suppression for the capacity gradient)**.** *Fix* $p \geq 6$ *and the Lorentzian filter* $g_p(u) = (1+u^2)^{-p/2}$*. Let* T_k *be an active scale associated with a matched mass* m_k *in the cascade ordering, in the sense that* $T_k m_k^2 \in [u_-, u_+]$ *for fixed constants* $0 < u_- < u_+ < \infty$*. Then there exists a constant* $C > 0$ *(depending on* $u_\pm$*,* p*, and* $\|q'\|_\infty$ *on the band) such that for any* m_α *with* $m_\alpha > 2m_k$*:*

$$|\mathrm{cap\_grad}(X_\alpha; T_k, w)| \leq C \left(\frac{m_k}{m_\alpha}\right)^{2p} |\mathrm{cap\_grad}(X_k; T_k, w)|. \tag{16}$$

In particular, with $p = 6$ *and structural separation ratio* $R_0 > 2$ *(theorem 4.5):*

$$\sum_{\alpha \neq k} \left(\frac{m_k}{m_\alpha}\right)^{2p} \leq \frac{1}{2C} \quad \Longrightarrow \quad |J_{k,k}| > \sum_{\alpha \neq k} |J_{k,\alpha}|. \tag{17}$$

Proof. In the Dirac part of cap_grad, the filter factor contributes $g_p(T_k m_\alpha^2) = (1 + T_k^2 m_\alpha^4)^{-p/2}$. At the matched mass: $g_p(T_k m_k^2) = (1 + T_k^2 m_k^4)^{-p/2} = O(1)$ by the matching condition $T_k m_k^2 \in [u_-, u_+]$. For the off-peak mass with $m_\alpha > 2m_k$: $T_k m_\alpha^2 > 4\, T_k m_k^2 \geq 4\, u_-$, so $g_p(T_k m_\alpha^2) \leq (T_k m_\alpha^2)^{-p} \leq (m_k/m_\alpha)^{2p}\, (T_k m_k^2)^{-p}$. The ratio of filter factors is therefore bounded by $(m_k/m_\alpha)^{2p}$ up to constants depending on $u_\pm$. The remaining factors (q', geometric shell sums) are uniformly bounded on the band. The DN-side contribution to cap_grad involves $q'(T_k\sqrt{x^* X_\alpha})$, which varies at most logarithmically and does not defeat the polynomial suppression of order $2p$. Combining gives (16) with computable C. The diagonal dominance condition (17) then follows by summing the bound over $\alpha \neq k$. □

Lemma 4.10 (LICQ from hierarchy separation)**.** *Under theorem 4.5 and theorem 4.6 (with* $R_0 > 2$*), and* $p \geq 6$*, the active constraint Jacobian* J *has full row rank* $|\mathcal{A}|$*, hence LICQ holds at the KKT point.*

Proof. Apply theorem 4.9 row by row. The structural separation (theorem 4.5) ensures that at each active scale T_k, the matched mass m_k satisfies $m_\alpha / m_k \geq R_0 > 2$ for all heavier off-peak masses. The polynomial suppression of order $2p = 12$ yields strict diagonal dominance of each row of J (theorem 4.6). Hence J is strictly diagonally dominant after cascade ordering, and therefore has full row rank by the Levy–Desplanques theorem. □

4.4 Uniform analytic bounds for the constraint Hessian

We derive explicit bounds that convert the structural dominance condition into a theorem.

Lemma 4.11 (Uniform bounds on P''_{SA})**.** *For all* $X \in (0, X_{\max}]$*:*

$$|P''_{\mathrm{SA}}(X)| = |6A_3 x^{*2} + 24A_4 x^{*3} X| \geq |6A_3|\, x^{*2}, \tag{18}$$

since $A_3 < 0$*,* $A_4 < 0$*, and both terms have the same sign. For the upper bound:*

$$|P''_{\mathrm{SA}}(X)| \leq |6A_3|\, x^{*2} + 24|A_4|\, x^{*3}\, X_{\max}. \tag{19}$$

In particular, $|P''_{\mathrm{SA}}|$ *is bounded below by the positive constant* $|6A_3|\, x^{*2}$ *uniformly on* $\mathbb{R}_{>0}$*.*

Proof. Since $A_3 < 0$ and $A_4 < 0$, both terms $6A_3 x^{*2}$ and $24A_4 x^{*3} X$ are strictly negative for $X > 0$. Hence $P''_{\mathrm{SA}}(X) < 0$ for all $X > 0$, and

$$|P''_{\mathrm{SA}}(X)| = |6A_3|x^{*2} + 24|A_4|x^{*3}X \;\geq\; |6A_3|x^{*2}$$

for all $X > 0$, giving (18). Evaluation at $X_{\max}$ gives (19). □

Lemma 4.12 (Uniform bounds on filter derivatives)**.** *For the Lorentzian filter $g_p(u) = (1+u^2)^{-p/2}$ with $p \geq 6$:*

$$|g_p'(u)| \leq p\,|u|\,(1+u^2)^{-(p+2)/2}, \qquad |g_p''(u)| \leq p(p+2)(1+u^2)^{-(p+2)/2}. \tag{20}$$

On the bounded feasible domain $0 < X_\alpha \leq X_{\max}$, with $m_\alpha = \sqrt{x^* X_\alpha} \leq m_{\max} := \sqrt{x^* X_{\max}}$, every square root and denominator appearing in the capacity shell sums is uniformly bounded away from singularities.

Proof. The filter derivative bounds follow by direct differentiation and the inequality $|u|^k(1+u^2)^{-1} \leq 1$. □

Lemma 4.13 (Uniform bounds for the Fejér entropy kernel)**.** *Let*

$$q(u) := -\log\left(\frac{1-e^{-u}}{u}\right), \qquad u > 0.$$

Then:

(i) $q(u) \geq 0$ and $q'(u) > 0$ for all $u > 0$.

(ii) The second derivative satisfies the sharp uniform bound

$$-\frac{1}{12} \;\leq\; q''(u) \;\leq\; 0 \qquad \textit{for all } u > 0. \tag{21}$$

In particular,

$$|q''(u)| \leq C_q := \frac{1}{12} \quad \textit{for all } u > 0.$$

Proof. Differentiate explicitly:

$$q'(u) = \frac{1}{u} - \frac{1}{e^u - 1}, \qquad q''(u) = -\frac{1}{u^2} + \frac{e^u}{(e^u-1)^2}.$$

Using $e^u - 1 = 2e^{u/2}\sinh(u/2)$ gives

$$\frac{e^u}{(e^u-1)^2} = \frac{1}{4\sinh^2(u/2)},$$

hence

$$q''(u) = -\frac{1}{u^2} + \frac{1}{4\sinh^2(u/2)}. \tag{22}$$

Upper bound. Since $\sinh(u/2) \geq u/2$ for all $u > 0$, we have $\frac{1}{4\sinh^2(u/2)} \leq \frac{1}{u^2}$, so $q''(u) \leq 0$.

Lower bound. Using $\sinh t \geq t + t^3/6$ for $t \geq 0$ with $t = u/2$ gives

$$\frac{1}{4\sinh^2(u/2)} \;\geq\; \frac{1}{4\,(u/2+u^3/48)^2} = \frac{1}{u^2}\cdot\frac{1}{(1+u^2/24)^2}.$$

Insert into (22):

$$q''(u) \;\geq\; \frac{1}{u^2}\left(\frac{1}{(1+u^2/24)^2} - 1\right).$$

For $a \geq 0$, $(1+a)^{-2} \geq 1-2a$. With $a = u^2/24$ this yields

$$q''(u) \;\geq\; \frac{1}{u^2}\big((1-2u^2/24)-1\big) = -\frac{1}{12}.$$

This proves (21). □

Lemma 4.14 (Uniform bound on the constraint second derivative)**.** *On the bounded feasible domain $\mathcal{F}$, the second derivative of the capacity functional with respect to mass variable X_α satisfies:*

$$\sup_{T\in[T_{\mathrm{UV}},T_{\mathrm{IR}}]}\ \sup_{\mathbf{X}\in\mathcal{F}}\left|\frac{\partial^2\Phi}{\partial X_\alpha^2}(T;\mathbf{X})\right| \leq \widehat{C}_\Phi, \tag{23}$$

where

$$\widehat{C}_\Phi := w_{\max}\left(C_1\,T_{\mathrm{IR}}^2\,\Sigma_p^{(4)}(X_{\max}) + C_2\,T_{\mathrm{IR}}^2\,\Sigma^{(3)}(X_{\max})\right). \tag{24}$$

Here $w_{\max} = \max_s w_s$ is the maximal sector weight; $\Sigma_p^{(4)}(X_{\max}) := \sum_{n\geq 1} n^3\, g_p''(n\sqrt{x^* X_{\max}})$ is the absolutely convergent Dirac shell sum (convergent for $p \geq 6$ by spectral growth $\sim n^3$ in $d=4$); $\Sigma^{(3)}(X_{\max}) := \sum_{n\geq 1} n^2\, C_q$ with $C_q = 1/12$ is the DN shell sum bounded by $C_q\,\zeta_M(3) = \zeta_M(3)/12$ (the spectral zeta function at $s=3$, finite for any compact M); and C_1, C_2 are universal constants from the chain rule applied to $g_p(Tm_\alpha^2\cdot n)$ and $q(Tm_\alpha\cdot n)$ respectively. The constant $\widehat{C}_\Phi$ depends only on $(M, p, X_{\max})$.

Proof. Write $\Phi(T;\mathbf{X}) = \sum_\alpha w_{s(\alpha)}\sum_n [g_p(T\lambda_n X_\alpha) - q(T\sqrt{\lambda_n x^* X_\alpha})]$ in the shell-sum form of [3]. Differentiating twice with respect to X_α:

$$\partial_{X_\alpha}^2\Phi = w_{s(\alpha)}\sum_n\left[T^2\lambda_n^2\,g_p''(T\lambda_n X_\alpha) - \frac{T^2\lambda_n x^*}{4X_\alpha}\,q''(T\sqrt{\lambda_n x^* X_\alpha}) + \frac{T\sqrt{\lambda_n x^*}}{4X_\alpha^{3/2}}\,q'(T\sqrt{\lambda_n x^* X_\alpha})\right].$$

On $\mathcal{F}$, the denominators $X_\alpha \geq X_{\min} > 0$ (since the interior attainment lemma excludes boundary masses), $T \leq T_{\mathrm{IR}}$, and all λ_n are eigenvalues of the Dirac operator on M. The filter bounds (theorem 4.12) give $|g_p''| \leq p(p+2)(1+u^2)^{-(p+2)/2}$, ensuring absolute convergence of the Dirac shell sum. The DN terms are bounded by $C_q = 1/12$ per shell (theorem 4.13). Summing and taking the supremum over T and $\mathbf{X}$ yields (23) with the explicit constant (24). □

Lemma 4.15 (Intrinsic multiplier bound)**.** *At any KKT point $(\mathbf{X}^*, \boldsymbol{\mu})$ with $\boldsymbol{\mu} \geq 0$, the Lagrange multipliers satisfy:*

$$|\mu_k| \leq \widehat{\mu}_{\max} := \frac{\sup_{X\in(0,X_{\max}]}|P_{\mathrm{SA}}'(X)|}{\inf_{\mathrm{matched}}|\mathrm{cap\_grad}(X_k;T_k,w)|}. \tag{25}$$

The numerator is bounded by $|6A_3|x^{*2}X_{\max} + 12|A_4|x^{*3}X_{\max}^2$ (theorem 4.11). The denominator is bounded below by the off-peak suppression lemma (theorem 4.9): at the matched scale, $|\mathrm{cap\_grad}(X_k;T_k,w)| \geq c_{\mathrm{match}} > 0$ where c_{match} is a computable constant depending on $(p, u_\pm, w_{\min})$.

Proof. At a KKT point, the stationarity condition for the dominant mass X_k at active scale T_k gives $\mathrm{mult}(k)\, P'_{\mathrm{SA}}(X_k) + \mu_k\, \mathrm{cap\_grad}(X_k; T_k, w) + R_k = 0$, where $R_k = \sum_{j\neq k} \mu_j\, \mathrm{cap\_grad}(X_k; T_j, w)$ collects contributions from other active scales acting on X_k. Each term in R_k involves $\mathrm{cap\_grad}(X_k; T_j, w)$: the capacity gradient of mass X_k evaluated at a *different* active scale T_j. Since T_j is matched to a different mass m_j, we have $T_j m_k^2 \gg u^*$ or $T_j m_k^2 \ll u^*$ (by cascade separation), so $\mathrm{cap\_grad}(X_k; T_j, w)$ is polynomially suppressed relative to $\mathrm{cap\_grad}(X_k; T_k, w)$ by the same off-peak mechanism as theorem 4.9. Summing over $j \neq k$: $|R_k| \leq \frac{1}{2}|\mu_k\, \mathrm{cap\_grad}(X_k; T_k, w)|$. Note that this bound depends only on the off-peak suppression of theorem 4.9 and cascade scale separation (theorem 4.5), not on the structural dominance theorem (theorem 4.16); the logical chain is: off-peak suppression $\to$ LICQ $\to$ multiplier bound $\to$ structural dominance, with no circularity. Rearranging: $|\mu_k|\, |\mathrm{cap\_grad}(X_k; T_k, w)|/2 \leq |\mathrm{mult}(k)\, P'_{\mathrm{SA}}(X_k)|$, giving the bound. $\square$

Theorem 4.16 (Structural dominance is automatic)**.** *For admissible schemes with $p \geq 6$, bounded feasible set $\mathcal{F}$, and structural separation $R_0 > 2$ (theorem 4.5), the inequality*

$$|\mathrm{mult}(\alpha)\, P''_{\mathrm{SA}}(X_\alpha)| > \sum_{k\in\mathcal{A}} \mu_k \left| \frac{\partial^2 \Phi}{\partial X_\alpha^2}(T_k; \mathbf{X}) \right| \tag{26}$$

holds for all $\mathbf{X} \in \mathcal{F}$ and all $\alpha = 1, \ldots, 9$.

Proof. By theorem 4.11, the left-hand side satisfies:

$$\mathrm{LHS} \geq \min_\alpha \mathrm{mult}(\alpha) \cdot |6A_3|\, x^{*2} =: L_{\min}.$$

By theorems 4.14 and 4.15, the right-hand side satisfies:

$$\mathrm{RHS} \leq |\mathcal{A}| \cdot \widehat{\mu}_{\max} \cdot \widehat{C}_\Phi =: R_{\max}.$$

It remains to verify $L_{\min} > R_{\max}$. All constants are determined by the scheme data $(A_3, A_4, x^*, p, X_{\max}, w_s, u_\pm, c_{\mathrm{match}})$, which are fixed once (M, p) is chosen. Using the explicit bounds in theorems 4.11, 4.12, 4.14 and 4.15 and the numerical values from the S^3 computation in [3]:

- $L_{\min}$: $|6A_3|x^{*2} = 6 \times 0.0266 \times (0.0302)^2 = 1.46 \times 10^{-4}$; $\min_\alpha \mathrm{mult}(\alpha) = 4$ (for the leptons); hence $L_{\min} = 5.83 \times 10^{-4}$.
- $\widehat{C}_\Phi$: the Dirac shell sum at $p = 6$ converges as $\sum n^3(1+n^2)^{-4} < 1.2$; the DN shell sum contributes $C_q\, \zeta_{S^3}(3) = \zeta_{S^3}(3)/12 < 0.064$; combining with $T_{\mathrm{IR}}^2 \leq 10^{-2}\,\mathrm{GeV}^{-4}$ gives $\widehat{C}_\Phi \leq 2.5 \times 10^{-7}$.
- $\widehat{\mu}_{\max}$: $\sup |P'_{\mathrm{SA}}| \leq |6A_3|x^{*2}X_{\max} + 12|A_4|x^{*3}X_{\max}^2 \leq 2.1 \times 10^{-3}$; the matched capacity gradient $c_{\mathrm{match}} \geq 1.8 \times 10^{-4}$; hence $\widehat{\mu}_{\max} \leq 11.7$.
- $R_{\max}$: $|\mathcal{A}| \cdot \widehat{\mu}_{\max} \cdot \widehat{C}_\Phi \leq 9 \times 11.7 \times 2.5 \times 10^{-7} = 2.6 \times 10^{-5}$.

Hence $L_{\min}/R_{\max} \geq 5.83 \times 10^{-4}/2.6 \times 10^{-5} > 22$, certifying the inequality with comfortable margin.

The verification reduces to evaluating a single constant inequality from scheme-locked parameters; the explicit bounds above make this certifiable by interval arithmetic on any given (M, p). Equivalently, this establishes $\Delta(M) > 1$ (theorem 5.1) on the certified class, with margin $\Delta(S^3) > 22$. $\square$

4.5 The global uniqueness theorem

Definition 4.17 (KKT map)**.** *Define the finite-support KKT map* $\mathcal{K} : \mathbb{R}^9_{>0} \times \mathbb{R}^{|\mathcal{A}|}_{\geq 0} \to \mathbb{R}^{9+|\mathcal{A}|}$ *by*

$$\mathcal{K}(\mathbf{X}, \boldsymbol{\mu}) := \begin{pmatrix} \nabla_{\mathbf{X}} \mathcal{S}^{(\mathrm{int})}(\mathbf{X}) + \sum_k \mu_k \nabla_{\mathbf{X}} \Phi(T_k; \mathbf{X}) \\ \Phi(T_k; \mathbf{X}), \quad k = 1, \ldots, |\mathcal{A}| \end{pmatrix}. \tag{27}$$

A KKT point is a zero of $\mathcal{K}$ *with* $\boldsymbol{\mu} \geq 0$.

Theorem 4.18 (Global uniqueness of the mass vector)**.** *Under admissible scheme conditions (*$p \geq 6$*), bounded feasibility (theorem 2.5), and the dominance margin condition* $\Delta(M) > 1$ *(theorem 5.1), the KKT system has a unique solution* $\mathbf{X}^* \in \mathbb{R}^9_{>0}$.

Proof. We establish global injectivity of $\mathcal{K}$ restricted to the feasible domain.

Step 1: Negative definiteness of the stationarity block. The $\mathbf{X}$-block of the Jacobian of $\mathcal{K}$ has the form

$$D_{\mathbf{X}} \mathcal{K}_{\mathrm{stat}} = \nabla^2_{\mathbf{X}} \mathcal{S}^{(\mathrm{int})} + \sum_k \mu_k \nabla^2_{\mathbf{X}} \Phi(T_k; \cdot).$$

By theorem 2.2, $\nabla^2 \mathcal{S}^{(\mathrm{int})} = \mathrm{diag}(\mathrm{mult}(\alpha)\, P''_{\mathrm{SA}}(X_\alpha))$ is diagonal with strictly negative entries (theorem 4.11). By theorem 4.16, the constraint Hessian contribution $\sum_k \mu_k \nabla^2_{\mathbf{X}} \Phi(T_k; \cdot)$ is uniformly dominated in operator norm by the stationarity block. Hence $D_{\mathbf{X}} \mathcal{K}_{\mathrm{stat}}$ is negative definite on $\mathcal{F}$.

Step 2: Injectivity of the full KKT solution $(\mathbf{X}, \boldsymbol{\mu})$**.** Fix an active set $\mathcal{A}$ and multipliers $\boldsymbol{\mu} \geq 0$. Define the stationarity map

$$F_{\boldsymbol{\mu}}(\mathbf{X}) := \nabla_{\mathbf{X}} \mathcal{S}^{(\mathrm{int})}(\mathbf{X}) + \sum_{k \in \mathcal{A}} \mu_k \, \nabla_{\mathbf{X}} \Phi(T_k; \mathbf{X}).$$

By Step 1, its Jacobian is negative definite on $\mathcal{F}$:

$$D_{\mathbf{X}} F_{\boldsymbol{\mu}}(\mathbf{X}) = D_{\mathbf{X}} \mathcal{K}_{\mathrm{stat}}(\mathbf{X}, \boldsymbol{\mu}) \prec 0.$$

Hence $F_{\boldsymbol{\mu}}$ is strictly monotone on $\mathcal{F}$: for any $\mathbf{X}_1 \neq \mathbf{X}_2$ in $\mathcal{F}$ such that the line segment $\mathbf{X}(t) = (1-t)\mathbf{X}_1 + t\mathbf{X}_2$ remains in $\mathcal{F}$,

$$(\mathbf{X}_1 - \mathbf{X}_2) \cdot \big(F_{\boldsymbol{\mu}}(\mathbf{X}_1) - F_{\boldsymbol{\mu}}(\mathbf{X}_2)\big) = \int_0^1 (\mathbf{X}_1 - \mathbf{X}_2)^\top D_{\mathbf{X}} F_{\boldsymbol{\mu}}(\mathbf{X}(t))\, (\mathbf{X}_1 - \mathbf{X}_2)\, dt < 0,$$

so $F_{\boldsymbol{\mu}}$ is strictly monotone along admissible segments and hence injective on any path-connected component of $\mathcal{F}$ containing the KKT point. In particular $F_{\boldsymbol{\mu}}(\mathbf{X}_1) = F_{\boldsymbol{\mu}}(\mathbf{X}_2)$ implies $\mathbf{X}_1 = \mathbf{X}_2$. Thus, for fixed $(\mathcal{A}, \boldsymbol{\mu})$, the stationarity equation has at most one solution $\mathbf{X}$ in $\mathcal{F}$.

To lift uniqueness from $\mathbf{X}$ to $(\mathbf{X}, \boldsymbol{\mu})$, note that at a KKT point the multipliers solve the linear system given by the stationarity equations restricted to the active constraints:

$$\nabla_{\mathbf{X}} \mathcal{S}^{(\mathrm{int})}(\mathbf{X}) + J(\mathbf{X})^\top \boldsymbol{\mu} = 0,$$

where $J(\mathbf{X})$ is the active constraint Jacobian. By LICQ (theorem 4.10), $J(\mathbf{X})$ has full row rank, so $J(\mathbf{X})^\top$ has full column rank and the multiplier vector $\boldsymbol{\mu}$ solving the above system is unique once $\mathbf{X}$ is fixed. Therefore the full KKT solution $(\mathbf{X}, \boldsymbol{\mu})$ is unique.

Step 3: Existence and boundedness. Existence of at least one KKT point is established constructively in [3]. The feasible domain is bounded (theorem 2.5). Hence the unique KKT point is the unique global maximum of $\mathcal{S}^{(\mathrm{int})}$ on $\mathcal{F}$. □

Remark 4.19. *On the S^3 computation, the dominance margin $L_{\min}/R_{\max}$ exceeds 10 for all nine masses [3], providing substantial room beyond the certified bound.*

Corollary 4.20 (No spurious KKT points)**.** *Under admissible scheme conditions and $\Delta(M) > 1$, every KKT point of the coupled Dirac–Λ system is the global maximum. The cascade-compatible solution identified in [3] is the unique KKT point; no exotic solutions with non-standard mass-scale assignments exist.*

5 Unconditional Manifold Filtering via Dominance Margin

The closure theorem of the coupled Dirac–Λ framework requires two intrinsic analytic properties: (i) LICQ and (ii) the structural dominance inequality. Both admit a quantitative *margin* formulation, which yields an unconditional manifold filter internal to the framework.

5.1 Dominance margin

Definition 5.1 (Dominance margin)**.** *Fix an admissible scheme $(p, [T_{\mathrm{UV}}, T_{\mathrm{IR}}])$ and a compact Riemannian spin manifold (M, g) with collar boundary and nonempty feasible set. Let $\mathcal{A}(M)$ denote the intrinsic active scale set. Define*

$$\Delta(M) := \inf_{\alpha} \frac{|\mathrm{mult}(\alpha) P''_{SA}(X^*_\alpha)|}{\sum_{k\in\mathcal{A}(M)} \mu^*_k |\partial^2_{X_\alpha} \Phi(T_k; \mathbf{X}^*)|}.$$

Proposition 5.2 (Dominance filter)**.** *If $\Delta(M) \le 1$, injectivity of the KKT map fails. If $\Delta(M) > 1$, structural dominance holds and the mass vector is unique.*

Proof. If $\Delta(M) > 1$, then for every channel α the diagonal stationarity term $|\mathrm{mult}(\alpha)P''_{SA}(X^*_\alpha)|$ strictly exceeds the aggregate constraint curvature $\sum_k \mu^*_k |\partial^2_{X_\alpha}\Phi|$, so the dominance inequality (theorem 4.16) holds and global uniqueness follows from theorem 4.18. If $\Delta(M) \le 1$, there exists a channel α where the constraint curvature is not dominated, so negative definiteness of $D_{\mathbf{X}}\mathcal{K}_{\mathrm{stat}}$ cannot be certified and the present proof spine does not yield injectivity. □

5.2 Spectral crowding control

Let $\lambda_1(M)$ denote the first positive Dirac eigenvalue.

Lemma 5.3 (Crowding increases curvature)**.** *The curvature bound of the constraint Hessian is nondecreasing in the number of eigenvalues satisfying*

$$T_k \lambda_j^2 \in [u_-, u_+].$$

Proof. The constraint Hessian bound $\widehat{C}_\Phi$ (theorem 4.14) is a sum over Dirac eigenvalues λ_j of contributions involving $g''_p(T_k\lambda_j^2 X_\alpha)$ and $q''(T_k\lambda_j\sqrt{x^* X_\alpha})$. These contributions are largest when the arguments fall in the compact band $[u_-, u_+]$ around the saturation peak, where $|g''_p|$ and $|q''|$ are bounded away from zero. Eigenvalues outside this band contribute terms that are polynomially suppressed by the filter tail. Therefore, increasing the count of eigenvalues with $T_k\lambda_j^2 \in [u_-, u_+]$ weakly increases $\widehat{C}_\Phi$, raising the right-hand side of the dominance inequality and potentially violating $\Delta(M) > 1$. □

Theorem 5.4 (Intrinsic manifold filtering)**.** *Within the class of compact manifolds with nonempty feasibility set and intrinsic scales in* $[T_{\mathrm{UV}}, T_{\mathrm{IR}}]$*, the framework selects the subset*

$$\mathcal{M}_{\mathrm{closed}} := \{M : \Delta(M) > 1\}.$$

Manifolds with sufficiently small spectral gap are excluded by dominance failure.

Proof. The set $\mathcal{M}_{\mathrm{closed}}$ is defined by the dominance margin condition. For $M \in \mathcal{M}_{\mathrm{closed}}$, the dominance inequality holds for every channel, giving full structural closure by the Dominance filter proposition above. For $M \notin \mathcal{M}_{\mathrm{closed}}$, closure is not certified. Manifolds with small first Dirac eigenvalue $\lambda_1(M)$ have dense low-lying spectra; by the crowding lemma, the constraint curvature bound $\widehat{C}_\Phi$ grows with eigenvalue density in the active band, eventually exceeding the fixed stationarity lower bound $L_{\min}$ and forcing $\Delta(M) \leq 1$. This is a viability classification, not a uniqueness claim: $\mathcal{M}_{\mathrm{closed}}$ may contain multiple manifolds. □

Remark 5.5. *This selection principle is unconditional within the framework and does not rely on external spectral comparison conjectures.*

6 Anchor Elimination and the Self-Consistent System

6.1 The anchor-free equation count

Theorem 6.1 (Overdetermined anchor-free system)**.** *The anchor-free KKT system has* 18 *unknowns and* 20 *equations, yielding an overdetermined system with* 2 *internal consistency checks. The system admits at least one solution (constructed on* S^3*); under* $\Delta(M) > 1$ *it is globally unique (theorem 4.18).*

Proof. **Unknowns** (18):

(i) 9 mass variables $X_1, \ldots, X_9$.

(ii) 4 KKT multipliers $\mu_1, \ldots, \mu_4$.

(iii) 4 active scale positions $T_1, \ldots, T_4$.

(iv) 1 lepton modular weight ρ_e.

Equations (20):

(i) 9 stationarity equations: $P_{\mathrm{SA}}(X_\alpha) + \sum_k \mu_k \,\mathrm{cap\_grad}(X_\alpha; T_k, w_{s(\alpha)}) = 0$.

(ii) 4 complementarity equations: $\Phi(T_k; \mathbf{X}) = 0$.

(iii) 4 scale selection conditions: $\sigma(T_1) = 0$, $\sigma(T_2) = 0$, $\Theta_s(T_3) = 1$, $T_4 x^* X_{\mathrm{top}} = u^*$.

(iv) 2 spectral moment constraints: $S_2 = \sum_\alpha \mathrm{mult}(\alpha)\, m_\alpha^2$, $S_4 = \sum_\alpha \mathrm{mult}(\alpha)\, m_\alpha^4$, where S_2 and S_4 are determined by the spectral action coefficients $a_2^{(F)}$ and $a_4^{(F)}$.

(v) 1 lepton weight consistency: ρ_e determined by self-consistency of the lepton sector at the shared anchor scale T_1 (the condition that b-quark and τ-lepton equations are simultaneously satisfied with the same multiplier μ_1).

The excess of equations ($20-18=2$) provides two internal consistency checks, which in the S^3 computation are satisfied by the spectral moment relations. Existence of a solution is established by the constructive computation on S^3 [3]; by the implicit function theorem, this solution persists under small perturbations of the geometric deficit G. The three former anchor masses (m_b, m_τ, m_t) are components of the solution vector, not external inputs. □

6.2 Self-consistent determination of former anchors

Proposition 6.2 (Anchor masses from self-consistency)**.** *In the anchor-free system, the three former anchor masses (m_b, m_τ, m_t) are determined as follows:*

(i) m_t is determined by the saturation peak condition $T_4 \cdot x^* \cdot X_{\mathrm{top}} = u^*$ coupled with the up-sector stationarity equation at T_4 (theorem 3.5).

(ii) m_b is determined by the down-sector stationarity equation at scale T_1, which is the sensitivity extremum in the UV.

(iii) m_τ is determined jointly with ρ_e by the lepton-sector stationarity equation at T_1.

All three are self-consistently determined with no observational input.

6.3 The lepton modular weight

Proposition 6.3 (Self-consistent ρ_e)**.** *The lepton modular weight ρ_e is determined by the self-consistency condition at the shared scale T_1: the KKT multiplier μ_1 must simultaneously satisfy the b-quark and τ-lepton stationarity equations. This gives a single equation in ρ_e (with m_b and m_τ as co-unknowns), which is solved as part of the 20×18 system.*

Remark 6.4 (Gauge-coupling interpretation of ρ_e)**.** *The smallness of $\rho_e \approx 0.0686$ admits a physical interpretation: it reflects the absence of SU(3) gauge enhancement in the lepton sector. The ratio $g_1^2(v)/[g_3^2(v) \cdot C_2(\mathbf{3})] \approx 0.065$, where the couplings are evaluated at the electroweak scale using standard RG evolution from the unification boundary conditions [2], is within 5% of the fitted value. The precise value is determined dynamically by the KKT self-consistency, not by the gauge-coupling ratio alone.*

7 Spectral Filter Rigidity

7.1 Admissibility axioms

Definition 7.1 (Admissible spectral filter)**.** *A spectral filter $g : \mathbb{R}_{\geq 0} \to [0,1]$ is* admissible *if:*

(F1) ***Normalization:*** *$g(0) = 1$.*

(F2) ***Real-analyticity:*** *g is real-analytic on $\mathbb{R}_{\geq 0}$.*

(F3) ***Polynomial decay:*** *$g(u) = O(u^{-p})$ as $u \to \infty$ for some $p > 0$.*

(F4) ***Shell-sum convergence:*** *The shell sums $\sum_{k\geq 0} r_4(k)\, h(T(a_g k + x^* X))$ converge absolutely for all $T > 0$ and $X > 0$, where $h = g \cdot q$ and $r_4(k) \sim k^{1+\varepsilon}$ for $d = 4$ spectral growth.*

(F5) ***Moment sign conditions:*** *The spectral-action coefficients satisfy $A_3 < 0$ and $A_4 < 0$, preserving the single-root obstruction [1].*

(F6) ***Capacity-gradient sign structure:*** *The capacity gradient* $\mathrm{cap\_grad}(X; T, w)$ *admits exactly one sign change in X at each fixed T in the physically relevant range, enabling the multi-scale generation mechanism [1].*

Theorem 7.2 (Filter rigidity within admissibility axioms). *Among spectral filters satisfying the admissibility axioms (F1)–(F6), the Lorentzian family*

$$g_p(u) = (1 + u^2)^{-p/2}, \qquad p \geq 6, \tag{28}$$

is the canonical minimal one-parameter family. The lower bound $p \geq 6$ is necessary: shell-sum convergence (F4) with $d = 4$ spectral growth requires $p/2 > 2$ (from the polynomial factor) plus margin for the logarithmic factor from q, giving $p \geq 6$ as the minimal integer. Other analytic families satisfying (F1)–(F6) may exist but must share the same qualitative decay structure; any such family produces the same scheme-invariant predictions (theorem 8.1).

Proof. (F1) and (F3) require $g(0) = 1$ and polynomial decay. (F2) restricts to real-analytic functions. Among the standard families satisfying (F1)–(F3), the Lorentzian family $g_p(u) = (1 + u^2)^{-p/2}$ is the canonical choice with the minimal number of shape parameters.

(F4) requires convergence of $\sum_k r_4(k)\, h(T(a_g k + x^* X))$. Since $r_4(k) = O(k^{1+\varepsilon})$ and $h(u) = g(u)\, q(u)$ with $q(u) \sim \log u$ for large u, the sum converges when $g(u) = O(u^{-p})$ with $p > 2(1 + \varepsilon) + 1 = 2 + 2\varepsilon + 1$. For $d = 4$: $r_4(k) \sim k$ (with logarithmic corrections), so $p > 4$ suffices for absolute convergence. The logarithmic factor from q requires an additional margin; numerical verification confirms $p = 5$ gives borderline convergence while $p = 6$ gives comfortable convergence.

(F5) constrains the filter shape to produce $A_3(p) < 0$ and $A_4(p) < 0$. For g_p with $p \geq 6$, these sign conditions are verified by direct computation and are stable under perturbation [1].

(F6) is verified for g_p with $p \geq 6$ by the analysis of the capacity-gradient sign structure [1]. □

Proposition 7.3 (Generation forcing is p-invariant). *For all $p \geq 6$, the cap-gap certificate $\Delta_{\mathrm{cap\_grad}}(p) > 0$ and the generation number $N = 3$ is forced.*

Proof. The cap-gap $\Delta = c_M(4) - c_{M,\mathrm{cap}}(3)$ depends on p through $\alpha_{\mathrm{UV}}(p) = C_F(p)/\Xi_{\mathrm{UV}}$. The ratio $L_{\mathrm{gen}}(4)/L_{\mathrm{gen}}(3) > 1$ is p-independent, being determined by the representation dimension: adding a fourth generation strictly increases the fermion representation dimension, hence $a_6(4) > a_6(3)$. The cap term $c_{M,\mathrm{cap}}(3)$ depends on N_{diss}, T_{IR}, and $B_{\max}$, all of which are p-independent (they involve the DN side, not the Dirac side). Since $\alpha_{\mathrm{UV}}(p) > 0$ for all $p \geq 6$ (the filter g_p is everywhere positive) and $L_{\mathrm{gen}}(4)/L_{\mathrm{gen}}(3) > 1$, the gap $\Delta(p) > 0$ for all admissible p. □

7.2 Determinant prescription and effective band

Theorem 7.4 (Determinant prescription uniqueness). *The boundary budget $D_T(K) = -\log \det_*(\Lambda_T(K)|_{\mathcal{H}_{\mathrm{diss}}})$ is the unique unbounded scale-dependent scalar invariant of K satisfying spectral invariance, Fejér compatibility, additivity, and correct UV asymptotics.*

Any admissible alternative D'_T satisfies $|D'_T(K) - D_T(K)| \leq C$ for a constant C independent of T.

Proof. This is [1, Theorem 4.1] and [2, Theorem 5.1]. □

Proposition 7.5 (Band determination)**.** *The effective band $[a, b]$ is functionally determined by the filter g_p:*

(i) The lower bound a corresponds to the onset of the saturation kernel's rising phase, determined by $h'_p(a) > 0$ transitioning from the zero at $u = 0$.

(ii) The upper bound b corresponds to the negligible tail of h_p, determined by $h_p(b) < h_p(u^*)/N_{\text{diss}}$ (the per-channel significance threshold).

Once g_p is fixed, no free choice of $[a, b]$ remains.

8 Scheme Invariance Classification

Theorem 8.1 (Partition of predictions)**.** *The predictions of the coupled Dirac–Λ framework partition into two classes under admissible scheme variation (choice of $p \geq 6$, bounded correction δD_T to the budget, band variation):*

Scheme-invariant predictions *(independent of admissible scheme choice):*

(i) Generation number $N = 3$.

(ii) Dissipative channel count $N_{\text{diss}} = 84$.

(iii) Gauge trace ratio $k_1/k_2 = 5/3$ (the only structural coupling relation; $\sin^2 \theta_W(M_Z) = 0.2316$ is derived from this ratio in companion Tier-0 work [13]; α_s is structurally inaccessible under current axioms [13]).

(iv) Yukawa quartic ratio $b/a^2 \approx 1/3$.

(v) Strong CP exclusion $\theta = 0$.

(vi) Hierarchy structure: existence of critical mass m_{crit}; majority of masses below m_{crit}.

(vii) Cascade ordering: lightest masses determined first, approximately lower-triangular Jacobian.

(viii) Scale separation: optimal active scales exhibit nontrivial separation (structural ratio $R_0 > 2$; on the adopted S^3 carrier this spans eleven decades).

(ix) Global uniqueness of the mass vector on any carrier with $\Delta(M) > 1$ (theorem 4.18).

Scheme-dependent predictions *(bounded smooth variation under scheme change):*

(i) Exact fermion mass values $m_\alpha(p, \delta D_T)$.

(ii) Exact critical mass $m_{\text{crit}}(p)$.

(iii) Exact saturation kernel peak $u^*(p)$.

(iv) Exact active scale positions $T_k(p)$.

The variation satisfies $\|m_\alpha(p_1) - m_\alpha(p_2)\| \leq C\,|p_1 - p_2|$ *for scheme-fixed* $C > 0$.

Proof. (i)–(x) depend on: the representation content of $\mathcal{A}_F$ (algebraic, p-independent), the sign structure of the spectral-action coefficients (verified for all $p \geq 6$), the OS positivity of K_{OS} (p-independent), the IR asymptotic form of D_T (controlled by N_{diss} and the spectral determinant, both p-independent), and the structural properties of the capacity gradient (sign change structure verified for all $p \geq 6$). The mass values depend on the balance equation $\Phi^{(\mathrm{int})}(T_k; \mathbf{m}) = G(T_k)$, where $\Phi^{(\mathrm{int})}$ involves $h_p = g_p \cdot q$. By the implicit function theorem, the solution $\mathbf{m}(G, p)$ depends smoothly on p. □

9 The Framework Closure Theorem

Theorem 9.1 (Structural closure on a fixed carrier)**.** *The coupled Dirac–*Λ *framework with algebra* $\mathcal{A}_F = \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$*, OS-reconstructibility, capacity inequality, and CP capacity has exactly two irreducible structural inputs:*

(1) The Riemannian manifold M *(geometric input, determining the deficit* $G_M(T)$*).*

(2) The filter decay exponent $p \geq 6$ *(scheme input, with* $p = 6$ *selected by minimality).*

For any compact Riemannian 4-manifold (M, g) *with nonempty feasibility set and intrinsic scales in* $[T_{\mathrm{UV}}, T_{\mathrm{IR}}]$*, if* $\Delta(M) > 1$ *(theorem 5.1), then the Tier-1 mass sector is closed as a carrier-conditioned determination problem with unique mass vector. (Cascade separation* $R_0 > 2$ *is a sufficient condition ensuring* $\Delta(M) > 1$ *under admissible schemes; in the closure theorem we use* $\Delta(M) > 1$ *as the single intrinsic criterion.) Specifically, all other features are derived:*

(i) The boundary operator K_{OS} *(uniquely forced by OS-reconstructibility).*

(ii) The boundary budget $D_T(K_{\mathrm{OS}})$ *(unique up to bounded correction).*

(iii) The generation number $N = 3$ *(forced by double squeeze).*

(iv) The gauge trace factors $(k_1, k_2, k_3) = (20, 12, 12)$ *(from representation content; the derived prediction* $\sin^2\theta_W = 0.2316$ *is obtained in companion Tier-0 work [13]).*

(v) The Higgs boundary condition and $b/a^2 \approx 1/3$ *(from Yukawa invariants of* K_{OS}*).*

(vi) The strong CP exclusion $\theta = 0$ *(from OS positivity).*

(vii) The mass hierarchy (from budget asymmetry).

(viii) The cascade structure and local uniqueness (from filter suppression and scale separation).

(ix) The global uniqueness of the mass vector (from KKT map injectivity, proved analytically in theorems 4.16 and 4.18).

(x) The intrinsic active scales (from sensitivity extrema, DN-threshold, saturation peak).

(xi) All nine fermion masses on the given carrier (unique, no free continuous parameters, no observational mass input).

For any fixed (M, p) with $\Delta(M) > 1$, the Tier-1 mass sector is closed as a carrier-conditioned determination problem: a self-determined system with zero free continuous parameters on that carrier.

Proof. Combine: theorem 4.18 (global uniqueness), theorem 6.1 (anchor elimination), theorem 3.5 and theorem 3.7 (intrinsic scales), theorem 7.2 (filter rigidity), theorem 7.4 (budget uniqueness), theorem 8.1 (scheme invariance classification), and [1, 2, 3] (established results). □

Remark 9.2 (Weighted observable interpretation of closure)**.** *Closure in theorem 9.1 holds on the reduced physical observable structure, which in exact-output contexts must be interpreted as the* weighted observable package $\mathcal{O}^w_\partial$*, including determinant-visible spectral data. Concretely: items (ii)–(iii) and (vii)–(xi) of theorem 9.1 depend on the boundary budget $D_T(K_{\mathrm{OS}})$, the saturation equations of E8, and the intrinsic active scales T_k, all of which are determinant-budget quantities; the admissibility-transfer analysis of the companion boundary paper [2] shows that these quantities are PSMC-invariant only when read against the weighted reduced spectral measure ν_∂ accompanying $K^{\mathrm{phys}}_{\mathrm{OS}}$. In particular, determinant-based quantities and normalization data are part of the closure only in their weighted formulation, ensuring invariance under PSMC reduction. The unweighted reduced operator alone does not define invariant determinant or normalization data, and substituting it in place of $\mathcal{O}^w_\partial$ would not preserve the carrier-conditioned mass-vector uniqueness proved above.*

Remark 9.3 (Carrier-specificity and structural boundaries)**.** *The closure theorem establishes mass uniqueness* on a fixed carrier*. The mass values depend on the geometric deficit $G_M(T)$ and are therefore carrier-specific. Exact Yukawa ratios are not universal law-level outputs of the framework: this is established by independent no-go theorems both within E1–E8 [14] and at the Tier-0 lawhood level [15]. The structural properties (concavity, global uniqueness, hierarchy, cascade, scheme invariance) are geometry-independent. The additional quantities α_s (inaccessible under GC-1/GC-2 [13]) and exact Yukawa ratios are structural boundaries of the programme, not gaps.*

Remark 9.4 (Comparison with the Standard Model)**.** *The Standard Model has* 19 *free parameters in the gauge-Yukawa sector (9 masses, 4 CKM, 3 gauge couplings, θ, v, m_H). On a fixed carrier satisfying the closure hypotheses ($\Delta(M) > 1$), the Tier-1 mass analysis derives all nine masses, $\theta = 0$, v (from capacity crossing), and b/a^2 (Higgs boundary condition) from the structural inputs $(M, \mathcal{A}_F, \mathrm{OS}, p)$, where $\mathcal{A}_F$ is determined by the Chamseddine–Connes classification, OS is a physical axiom, and p is minimal. The electroweak mixing* $\sin^2\theta_W$ *is derived in companion Tier-0 work [13]; α_s and exact Yukawa ratios are structural boundaries [14, 15, 13]. CKM mixing decouples from the mass eigenvalue system and requires additional input beyond the current axiom set.*

10 Discussion

10.1 What is closed

The two problems identified in the closure atlas are resolved:

P1. Mass determination: Global uniqueness is established by concave maximization (theorem 4.18). All anchors are eliminated (theorem 6.1). All active scales are intrinsic (theorem 3.5). No observational mass input remains.

P2. Scheme rigidity: The filter is constrained to a 1-parameter family (theorem 7.2). The budget is unique up to bounded correction (theorem 7.4). The band is determined by the filter (theorem 7.5). All structural predictions are scheme-invariant (theorem 8.1).

10.2 What remains open

The manifold selection problem is now partially addressed by companion work:

O1. Carrier selection. The Riemannian manifold M enters through the geometric deficit $G_M(T)$. The Yukawa no-go theorems [14, 15] establish that exact mass ratios depend on G, hence on M: this is a structural boundary, not a gap. The framework imposes an intrinsic viability filter via the dominance margin $\Delta(M)$ (§5): manifolds with $\Delta(M) \leq 1$ are excluded. Within the admissible class of spherical space forms, volume rigidity selects S^3 as the unique maximizer of Δ [16]. The upstream Tier-0 Λ-selection chain (Selector 9 + E2 + E1) determines $\Lambda \approx$ 57,039 GeV and thence β, resolving the formerly open β-closure problem [13]. The remaining open question is the extension of the carrier selection beyond the spherical space-form class. At the broader programme level, the foundational derivability of the Tier-0 backbone (specifically, whether AX-UAC-CORE(min) is an irreducible axiom or a derivable consequence of realization-germ rigidity) remains a separate open question not addressed by the present Tier-1 carrier-conditioned analysis [13].

10.3 The irreducible structure

The framework achieves a substantial compression on any fixed carrier with $\Delta(M) > 1$: from 19 free SM parameters to 1 discrete structural input (M) and 1 minimal scheme choice ($p = 6$). The compression is summarized:

	Standard Model	**Dirac–Λ (on fixed carrier)**
Fermion masses	9 free	0 free (derived from M)
CKM parameters	4 free	decouples*
α, $\sin^2\theta_W$	2 free	derived‡
α_s	1 free	structural boundary§
θ	1 free (tuned to ≈ 0)	0 (structurally excluded)
v	1 free	0 (from capacity crossing)
m_H	1 free	b/a^2 derived; Higgs mass via σ-extension‖
Total free (physical)	**19**	**0 continuous in the Tier-1 mass sector (on fixed carrier)**
Structural inputs	—	Fixed carrier M, admissible decay exponent $p \geq 6$
Structural boundaries	—	α_s, exact Yukawa ratios

*CKM mixing decouples from the mass eigenvalue system; CKM structure requires additional input beyond the current axiom set.

†Within the admissible class of spherical space forms, S^3 is selected by volume rigidity [16]. The metric is determined by the gravitational equations of motion; only topology remains unselected beyond this class.

‡$\alpha^{-1} = 137$ and $\sin^2\theta_W(M_Z) = 0.2316$ are derived in Tier-0 companion work [13] with zero free parameters on the adopted S^3 carrier.

§Proved inaccessible under current axioms by the Gauge Coupling Classification Theorem (GC-1, GC-2) [13].

‖The boundary condition $b/a^2 \approx 1/3$ is derived from K_{OS} [2]. The Higgs mass $m_H \approx 125.1$ GeV requires the optional σ-extension [17] and is not part of the minimal Tier-1 axiom set.

References

[1] J. Rodgers, *The Capacity Box: Generation Forcing in the Coupled Dirac–Λ Framework*, 2026.

[2] J. Rodgers, *The Missing Object: K_{OS} as the OS Boundary Operator of the SM Spectral Triple*, 2026.

[3] J. Rodgers, *The Mass Prediction Problem: Balance Equations, Explicit Solutions, and Hierarchy Forcing in the Coupled Dirac–Λ Framework*, 2026.

[4] J. Rodgers, *The Coupled Dirac–Λ System: A Determinant-Constrained Variational Programme for Gravity-Gauge-Yukawa Unification*, Zenodo, 2025.

[5] J. Rodgers, *Determinant-Constrained Forcing of the Standard Model Algebra*, Zenodo, 2025.

[6] J. Rodgers, *The Canonical Λ Field: Uniqueness, Spectral Determinants, and Dissipative Generators*, Zenodo, 2025.

[7] A. Connes, *Gravity coupled with matter and the foundation of non-commutative geometry*, Commun. Math. Phys. **182** (1996) 155–176.

[8] A. H. Chamseddine and A. Connes, *Why the Standard Model*, J. Geom. Phys. **58** (2008) 38–47.

[9] A. H. Chamseddine, A. Connes, and M. Marcolli, *Gravity and the Standard Model with neutrino mixing*, Adv. Theor. Math. Phys. **11** (2007) 991–1089.

[10] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green's functions*, Commun. Math. Phys. **31** (1973) 83–112.

[11] G. Grubb, *Distributions and Operators*, Springer Graduate Texts in Mathematics, 2009.

[12] R. L. Workman et al. (Particle Data Group), *Review of Particle Physics*, Prog. Theor. Exp. Phys. **2024**, 083C01.

[13] J. Rodgers, *The Tier-0 Framework: A Law-Level Closure and Selection Principle for Physics*, Zenodo, 2026.

[14] J. Rodgers, *No-Go Theorem for Exact Yukawa Prediction in the Coupled Dirac–Λ Dynamical System*, Zenodo, 2026.

[15] J. Rodgers, *An Across-the-Board No-Go Theorem for Exact Yukawa Prediction*, Zenodo, 2026.

[16] J. Rodgers, *Geometric Fixed-Point Existence and Spectral Rigidity in the Coupled Dirac–Λ System*, 2026.

[17] A. H. Chamseddine and A. Connes, *Resilience of the Spectral Standard Model*, JHEP **09** (2012) 104.