The First χ-Increment Magnitude-Bridge and Charm-Hinge Theorem in VERSF

▲ Programme Milestone — Flavour-Magnitude and Mass-Hierarchy Series Gate FH-1 / First χ Spectral Gap, Charm-Hinge Selection, Non-Insertion Yukawa Audit, and Rare-Decay Separation

Deriving the carrier of the first flavour-magnitude from the spectral structure of an ordered χ-access operator, selecting the charm hinge by substrate-defined discriminators rather than mass labels, and firewalling the result against charm-loop effects in rare neutral-current decays

Keith Taylor VERSF Theoretical Physics Programme Flavour-Magnitude and Mass-Hierarchy Series Tier A gate-theorem paper, FH-1

General Reader Summary

The up quark, charm quark, and top quark carry the same type of electric charge, yet their masses differ enormously. In the Standard Model those differences are encoded in Yukawa couplings: each quark receives its own number, set by hand to match experiment. That is successful bookkeeping, but from a first-principles standpoint it raises a serious question:

Did we derive the flavour hierarchy, or did we simply assign a separate fitted number to every particle?

The Higgs-radial theorem establishes that VERSF has a legitimate scalar mass carrier: the radial stiffness mode of the closure interface. But a carrier is not a hierarchy. Something must still explain why the up-type channels split into light, intermediate, and heavy support.

VERSF contains an ordered flavour-access structure, denoted χ. In this paper χ is sharpened from a verbal ordering into a physical spectral object: a positive operator on the space of up-type channels, whose eigenvalues are the amounts of access to mass-generating support actually realised by the three channels. Whether these are also the *lowest* amounts the substrate would permit — whether nothing sits unoccupied beneath the first step — is a named premise inherited for verification from the census gate, not a silent assumption. Three logically different things can then be established, and this paper is careful to keep them apart:

1. **Order** — which channel sits above which (ordinal closure);
2. **Gap** — the size of the first separation above the floor, convention-independent once the scale of the access operator is physically fixed — a fixing this paper names as an explicitly open item (metric closure);
3. **Number** — the actual Yukawa value that separation produces (dynamical closure).

Under the explicitly stated spectral, channel-count, floor, endpoint, alignment, and identification premises, this paper proves the first, defines exactly what the second requires, and constructs the architecture — the magnitude-bridge — through which any future claim to the third must pass. It does not compute the charm mass, and it says so in every place where a reader might otherwise assume it does.

The selection of charm is not made by pointing at the middle quark. The paper defines three substrate-level discriminators — floor proximity, interior isolation, and endpoint saturation — and shows that they select a unique interior, first-lifted channel *before any mass data or generation names are consulted*. Only then is the channel named, in two stages: gauge quantum numbers fix that it is an up-type quark — they cannot distinguish the three generations, which is the flavour puzzle itself — and the alignment between access channels and mass channels pairs the selected first-lifted channel with the intermediate-mass one, which is what "charm" means in the Standard Model. That second stage carries its own named condition: the non-access factors in the coupling — normalisation, carrier strength, dressing — must not be so unbalanced across channels that they overturn the access ordering. This is stated as a premise with its own theorem and falsifier, not assumed silently; if it failed, the paper would still have selected the first access step, but could no longer call it charm. Mass information enters only in the final naming convention, never in the selection. The identification is also shown to be stable: small admissible corrections cannot silently move the hinge to a different channel.

One clarification of vocabulary: charm's separate appearance as a loop contribution in rare $b \to s\, \mu^{+}\mu^{-}$ decays is a different typed object from its mass-ladder role, and the two may not be used as evidence for or against each other without a dedicated bridge (Theorem 8, §10″).

In one sentence:

Under the stated premises, the charm channel is the unique minimal admissible non-floor up-type channel within the three-state χ spectrum — the first metrically nonzero step of ordered flavour access — selected by substrate discriminators, identified by the two-stage bridge (P13), stable under sub-gap perturbations, and firewalled from rare-decay loop physics.

Scope and Status Box

This box merges the scope declaration and the predictive-content ledger. Status vocabulary: **Inherited** (supplied by a prior gate), **Premised** (assumed here with a named falsifier), **Conditionally proved** (proved from the named premises), **Closed** (audit structure fully supplied here), **Open** (explicit debt, see §10′).

Claim	Status
Higgs-radial carrier as prior mass-support structure	Inherited

Claim	Status
Flavour hierarchy requires structure beyond the Higgs-radial carrier	Closed
χ as an ordered flavour-access structure	Premised / inherited
χ as a positive self-adjoint spectral operator with well-founded spectrum	Premised (P6)
First positive χ increment (spectral gap $\delta\_\chi,_1$) exists	Conditionally proved from well-founded spectrum
$\delta\_\chi,_1$ carries an invariant magnitude (Level II)	Displacement architecture closed; calibration and value open (D2, D3)
Non-χ factors preserve the mass ordering	Premised (P8A); Theorem 6A conditional on it
Φ and Σ constructed from substrate operators	Open (D15); Γ spectrally constructed — Closed
Magnitude enters mass support only through the factorised bridge	Conditionally closed under P9 (projector commutation)
Factorisation grounded at operator level (block-diagonal effective action)	Conditionally proved (P9)
Floor channel excluded as clean witness (via discriminator Φ)	Conditionally proved
Endpoint channel excluded as first-increment witness (via discriminator Σ)	Conditionally proved
Unique interior hinge selected by (Φ, Γ, Σ) without mass data	Conditionally proved within three-state bridge
Selected hinge identified with charm via the two-stage bridge	Conditionally proved (P13, Theorem 6A)
Hinge identity stable under sub-gap perturbations	Conditionally proved (Theorem 9)
χ/mass alignment at FH-1 order	Premised (P17A); derivation Open (D13)
Channels occupy the lowest admissible strata of $\mathcal{X}_$phys	Premised (P6A, SC-2 inheritance); derivation Open (D14)
χ scale fixed only up to unitary equivalence; recalibration fixing	Open (D2)
Hinge witnesses the first admissible, not merely first occupied, increment	Conditionally proved (Lemma 1)
Non-insertion of the charm Yukawa enforced by parameter partition	Closed (Theorem 10)
Charm Yukawa target scheme-declared and dimensionless	Closed
QCD/confinement dressing separated from χ magnitude	Closed
Rare-decay/mass-hierarchy type separation	Closed (Theorem 8)

Claim	Status
Magnitude of $\delta\_\chi,_1$ physically derived	Open (D2, D3)
Access-to-mass law $\mathcal{M}\_\chi$ derived	Open (D4)
Charm Yukawa numerically derived	Open (D1–D7)
Up mass, top mass, full hierarchy	Open / outside scope
Down-type, lepton, neutrino analogues; CKM mixing	Outside scope
$b \to s\ \mu^+\mu^-$ effective-operator structure, Wilson-coefficient shifts	Outside scope (D12)
Empirical confirmation	Outside scope

FH-1 is therefore a **first-gap admissibility and magnitude-bridge theorem**. It conditionally closes the ordinal level, fully specifies the metric level, and constructs the audit architecture for the dynamical level, without claiming the dynamical level itself.

Abstract

The Higgs-radial gate establishes a positive-sector scalar mass carrier as the normal stiffness mode of the VERSF closure interface. The existence of a carrier does not derive the fermion mass hierarchy. A separate theorem must explain why ordered flavour access contains a first positive separation, what kind of object that separation is, and why the charm channel — rather than an arbitrary second-generation Yukawa insertion — is where it first becomes visible.

This paper proves the First χ-Increment Magnitude-Bridge and Charm-Hinge Theorem. The χ readout is sharpened into a positive, invariant, self-adjoint χ-access operator $\hat{\chi}$ on the physical up-type channel space $\mathcal{U}$. Under the premise that its physical spectrum is well founded — a lowest support value λ_0 exists and the set of strictly greater admissible values possesses a least element λ_1 — the first physical χ gap

$\delta\_\chi,_1 = \lambda_1 - \lambda_0 > 0$

is a defined spectral object rather than a verbal ladder step. Where χ is not globally affine, $\delta\_\chi,_1$ is the invariant spectral or geodesic distance between the lowest two admissible strata; the symbol $\Delta\_\chi(U\_F, U\_C)$ denotes the typed displacement, with $\delta\_\chi,_1 = \|\Delta\_\chi(U\_F, U\_C)\|\_\chi$ its invariant scalar magnitude, and coordinate subtraction is never assumed.

Three closure levels are distinguished and kept separate throughout: Level I (ordinal closure — an invariant ordering exists), Level II (metric closure — the first separation has an invariant nonzero magnitude), and Level III (dynamical closure — a substrate-derived map $\mathcal{M}\_\chi(\delta\_\chi,_1; \lambda\_\text{substrate})$ converts the separation into reduced mass support, with every parameter fixed independently of the observed charm mass). FH-1 conditionally closes Level I, specifies the requirements of Level II, and constructs the architecture of Level III without claiming it.

The hinge is selected without mass data. Three substrate-defined discriminators are introduced — floor proximity $\Phi(U)$, interior isolation $\Gamma(U)$, and endpoint saturation $\Sigma(U)$ — and the first clean hinge is defined by the selection rule

$$U_* = \operatorname{argmin} \{ d_\chi(\chi(U), \chi_0) > 0 : \Phi(U) < \Phi_{max} , \Sigma(U) < \Sigma_{max} , \Gamma(U) \geq \Gamma_{min} \} ,$$

taken over $U \in \mathcal{U}_{phys}$. Only after U_* is fixed does a two-stage bridge name it: representation content $(3, 2, +1/6)_L \oplus (3, 1, +2/3)_R$ fixes the up-type sector — it cannot select among the three generations, which carry the identical datum — and the P17A alignment together with the Theorem 6A (P8A) order transfer pairs the first-lifted χ channel with the intermediate-mass eigenchannel, which is by Standard Model definition the charm channel. Mass ordering enters exactly once, at that definitional naming step, never in the selection. An alignment premise (P17A) further guarantees that the χ eigenchannel and the mass eigenchannel coincide at FH-1 order, approximate alignment being admissible only under a declared spectral-projector (Davis–Kahan-type) bound deferred to D13, so the identification is basis-invariant; a stability theorem shows the hinge identity survives all admissible perturbations below half the smaller adjacent spectral gap. An occupancy premise — stated as the FH-1 restriction of the SC-2 world-to-framework census, its discharge by SC-2 an open verification (D14) — guarantees, where "admissible" is claimed, that the physical channels occupy the lowest admissible strata of the support set $\mathcal{X}_{phys}$; unconditionally the paper delivers the first occupied gap, and under P6A the first admissible one (Lemma 1).

The admissible charm-channel structure is the dimensionless factorised bridge

$$y_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathcal{R}_H(\mu) \cdot \mathcal{F}_\chi(\delta_{\chi,1}) \cdot \mathcal{D}_C(\mu) ,$$

grounded at operator level in a block-diagonal effective action with multiplicative renormalisation, so that the four scalar factors arise as the χ-eigenchannel projection of

$$Y_u(\mu) = Z_L^{(-1/2)}(\mu) \cdot B_u \cdot f_\chi(\hat{\chi}) \cdot Z_R^{(-1/2)}(\mu)$$

rather than being postulated. A parameter-partition theorem formalises non-insertion: the parameter set is partitioned into prior, FH-1, and empirical classes, the empirical class is required to be disjoint from all observed charm data during derivation, and a local identifiability rank condition ensures no normalisation constant can smuggle the charm value in through the back door.

A rare-decay separation theorem additionally types charm-hinge mass support and charm-loop transition support as distinct objects, with all cross-inference deferred to a dedicated effective-operator bridge (debt D12).

Eleven theorems and a lemma are proved.

Theorem 1 — χ-Gap Necessity Theorem. A nondegenerate, well-founded physical χ spectrum possesses a first positive gap.

Theorem 2 — First-Gap Magnitude Theorem. The first gap carries a convention-independent magnitude if and only if the χ operator's spectral calibration (or an equivalent invariant distance) is physically fixed; when it is, $\delta\_\chi{,}_1 > 0$ is defined independently of observed masses.

Theorem 3 — Mass-Bridge Factorisation Theorem. An admissible charm-channel Yukawa must arise as the χ-eigenchannel projection of a block-diagonal, multiplicatively renormalised effective Yukawa operator, yielding the four-factor bridge.

Theorem 4 — Light-Floor Exclusion Theorem. The floor-proximate channel (Φ ≥ Φ_max) cannot serve as the clean first-gap witness.

Theorem 5 — Endpoint Exclusion Theorem. The saturation channel (Σ ≥ Σ_max) cannot serve as the first-gap witness.

Theorem 6 — Hinge Selection and Charm Identification Theorem. The discriminator selection rule fixes a unique interior first-lifted channel U_* without mass data; the two-stage bridge identifies U_* with charm, conditional on Theorem 6A.

Theorem 6A — Net Mass-Order Preservation Theorem. If each adjacent log-f_χ gap exceeds the corresponding non-χ log-ratio (P8A), the full couplings satisfy $y_0 < y_1 < y_2$, so the first-lifted χ channel is paired with the intermediate-mass eigenchannel. Without it, χ order does not license mass order.

Theorem 7 — Charm Non-Insertion Theorem. Within the FH-1 audit class, the factorised, independently normalised, non-inserted χ bridge is necessary and gate-sufficient for admitting charm as the first-gap candidate. It is not sufficient for full phenomenological validity.

Theorem 8 — Rare-Decay/Mass-Hierarchy Separation Theorem. Charm-hinge support and charm-loop rare-decay support are distinct typed objects; no cross-inference is admissible absent the D12 bridge.

Theorem 9 — Hinge Stability Theorem. Under any admissible perturbation smaller than half the minimum adjacent gap, the eigenvalue ordering is unchanged and the first-lifted eigenchannel remains continuously connected to U_*.

Theorem 10 — No-Parameter-Hiding Theorem. Under the parameter partition, the charm-proxy exclusion with declared insensitivity bound, and the identifiability rank condition, no factor of the bridge can function as a hidden charm fit — direct or by proxy.

FH-1 closure grade: Level I conditionally closed under P2–P6, with P6A additionally where "admissible" is claimed; Level II displacement architecture closed, physical calibration and numerical value open (D2, D3); Level III open (D1–D7); charm identification conditional on net mass-order preservation (P8A, Theorem 6A) and exact alignment (P17A); rare-decay typing closed.

Contents

0. Inheritance, Aim, and Firewalls

0.1 Inherited VERSF disciplines

I1 — Physical-sector admissibility. A mass-channel claim is not admitted merely because a numerical value can be fitted. It must preserve positivity, invariance, and typed physical support.

I2 — Non-insertion discipline. Standard Model Yukawa values may not be copied into VERSF and renamed as derived χ data.

I3 — Higgs-radial inheritance discipline. Mass support must pass through the Higgs-radial carrier derived at the closure-interface gate.

I4 — Bridge discipline. A channel label is not physical unless a bridge is stated from the internal carrier to a downstream mass, Yukawa, spectral, or observable role.

I5 — Scheme discipline. Quark masses and Yukawa values must be stated in a declared renormalisation scheme and scale.

I6 — Non-double-counting discipline. Higgs-radial support, χ magnitude, sector normalisation, QCD running, and mixing effects may not be counted multiple times under different names.

I7 — Positivity discipline. χ structure feeding physical mass channels must be positive-sector, non-ghost, non-gauge, and invariant under permitted relabellings.

I8 — Ordering discipline. χ order may not be rearranged after seeing observed masses.

I9 — Gate-admissibility discipline. A gate is closed only when its primitives, readouts, normalisations, bridge debts, and falsifiers are explicit.

I10 — Observable-typing discipline. Distinct experimental observables belong to distinct audit classes and may not be pooled as a single diagnostic without a typing bridge.

0.2 The new problem and the core aim

The Higgs-radial theorem answers: why is there a scalar mass carrier at all? It does not answer: why do different fermion channels receive different amounts of mass support?

A first-principles derivation cannot write y_u, y_c, y_t as three independent fitted numbers and claim the up-type hierarchy has been derived. FH-1 asks the narrower prior question — where does the first nonzero up-type flavour separation appear, and what kind of object is it? — and has one aim:

Under the explicitly stated χ-spectrum, channel-count, floor, endpoint, alignment, and identification premises, to prove that a unique interior first-lifted channel exists, is selected by substrate discriminators without mass data, is identified with charm by a two-stage bridge (sector by representation content, channel by alignment and order transfer), is stable under sub-gap perturbations, and can carry Yukawa support only through a factorised, non-inserted magnitude-bridge — while remaining firewalled from charm-loop physics in rare neutral-current decays.

0.3 Firewalls

This paper does **not** claim: the numerical charm mass; the charm Yukawa without declared scheme and scale; the full up/charm/top hierarchy; the up or top masses; that QCD running or confinement effects are negligible; CKM mixing; down-type, lepton, or neutrino analogues; that $\delta\_\chi{,}_1$ may be chosen by fitting charm; that the Standard Model charm Yukawa may be imported and renamed as χ structure; that the χ metric or the access-to-mass law is derived here unless explicitly inherited; that the $b \to s\ \mu^+\mu^-$ tensions are evidence for FH-1; that FH-1 controls charm-loop amplitudes or Wilson-coefficient shifts; or that passing FH-1 is empirical confirmation.

---

1. The First-Gap Problem

1.1 Why charm is dangerous in a derivation

Charm is dangerous because it is tempting to treat it as "just the second generation." That phrase assumes what must be proved: that there is a ladder, that the ladder is ordered, that the second up-type channel is not arbitrary, that its mass is not a fitted parameter, that its relations to the Higgs-radial carrier, the light floor, and the top endpoint are known, that QCD running is accounted for, and that mixing is not being confused with mass support.

There is a second, sharper danger, and it is the one this paper confronts directly. A classification that *premises* a floor role, an endpoint role, and an intermediate role, then eliminates the floor and endpoint, has proved uniqueness only inside its own chosen classification. The referee question is legitimate:

"Where is the new physical derivation, as opposed to a labelled elimination?"

FH-1 answers it in §3 and Theorem 6: the three roles are defined by substrate-level discriminators Φ, Γ, Σ that make no reference to observed masses or generation names; the hinge is selected by an explicit rule from those discriminators alone; and only then does the two-stage bridge (P13) — not a counting convention — identify the selected channel with charm.

1.2 What must be derived

The first flavour-magnitude gate must derive or explicitly premise: why an ordered access structure exists; why it has a first positive gap; why the gap is an invariant magnitude rather than an ordinal label; through what structure the gap may enter mass support; why the floor and endpoint channels are excluded as clean witnesses by independent criteria; why the selected hinge is charm under the two-stage bridge (P13); why the identification is basis-invariant and perturbation-stable; why the charm Yukawa cannot be inserted by hand; and why the hinge role is separated from rare-decay loop physics. FH-1 addresses exactly these points and names which are premised rather than proved.

1.3 The core intuition

The up channel sits at the floor. The top channel sits near the ceiling. Charm is the first visible stair.

Charm is the first place where the up-type ladder rises high enough to show the χ gap, but not so high that endpoint saturation hides the first-step structure.

The rest of the paper converts that intuition into spectral, selectional, and stability statements.

---

2. Named Premises of the χ-Gap Audit

Premises are labelled P1–P19, with P17A the alignment premise.

2.1 P1 — Higgs-radial inheritance premise

VERSF has an admissible Higgs-radial mass carrier R_H, inherited from the closure-interface radial-stiffness gate.

Falsifier. FH-1 attempts to derive mass-channel support without any Higgs-radial carrier or equivalent mass-support bridge.

2.2 P2 — χ-access premise

There exists a VERSF flavour-access readout on the physical up-type channel space $\mathcal{U}$, realised at this gate as a χ-access operator (P6) whose spectrum supplies the ordered support values.

Falsifier. The paper claims χ structure while no readout, operator, or equivalent ordered access coordinate is defined.

2.3 P3 — χ-ordering premise

The χ-support values carry an invariant order relation $\preceq_\chi$. If $\chi_A \prec_\chi \chi_B$, then channel B has strictly higher admissible flavour-access support than channel A.

Falsifier. The channel order changes under labels, conventions, or observed mass fitting.

2.4 P4 — χ-independence premise

χ is not defined as observed mass. It is an independent access structure whose relation to mass must be bridged.

Falsifier. χ_C is defined by reading off the charm mass, and the charm mass is then claimed as derived from χ.

2.5 P5 — Nondegeneracy premise

The up-type sector is not fully χ-degenerate: at least two channels differ in admissible χ support.

Falsifier. All up-type channels have identical χ support while a hierarchy is still claimed.

2.6 P6 — Physical χ-spectral-gap premise

There exists a positive self-adjoint χ-access operator, restricted to the admissible up-type channel space by the physical projection,

$$\hat{\chi}_{phys} = P_{phys} \cdot \hat{\chi} \cdot P_{phys} ,$$

invariant under the permitted gauge, frame, and basis transformations.

Its physical spectrum contains a lowest support value λ_0, and the positive upper set

$$S_+(\lambda_0) = \{ \lambda \in \mathrm{Spec\_phys}(\hat{\chi}) : \lambda > \lambda_0 \}$$

is nonempty and possesses a least element λ_1. The first physical χ gap is therefore

$$\delta_{\chi,1} = \lambda_1 - \lambda_0 > 0 .$$

Where χ is not globally affine, $\delta\_\chi,_1$ is understood as the invariant spectral or geodesic distance between the lowest two admissible strata, never as coordinate subtraction.

The invariance class of $\hat{\chi}$ is unitary equivalence only. A monotone spectral recalibration $\hat{\chi} \to f(\hat{\chi})$ preserves the ordering, the eigenvectors, the discriminator structure, and every ordinal fact while changing every gap; such recalibrations are therefore not admissible transformations at this gate. Which member of the recalibration class is the substrate operator is a physical scale-fixing that belongs to the χ-metric debt D2, and Level II invariance is conditional on its discharge.

Falsifier. No positive physical χ operator, invariant support functional, or well-founded admissible spectrum is supplied; the alleged first gap depends on coordinate choice; or the positive spectrum accumulates at the floor so that no least positive separation exists; or the gap magnitude is treated as convention-independent while the monotone-recalibration freedom of $\hat{\chi}$ remains unfixed.

2.6′ P6A — Lowest-strata occupancy premise

The substrate-admissible support set $\mathcal{X}$_phys (§3.1) may exceed the realised spectrum. P6A states that the physical up-type channels occupy the lowest admissible strata:

Spec_phys($\hat{\chi}$) = the three $\preceq\_\chi$-least elements of $\mathcal{X}$_phys ;

in particular, no admissible stratum $\chi \in \mathcal{X}$_phys with $\chi_0 \prec\_\chi\ \chi \prec\_\chi\ \lambda_1$ exists unoccupied, and the floor stratum is the $\preceq\_\chi$-least admissible value. (Lemma 1 uses only these two clauses; the full lowest-three form is retained, strictly stronger, for downstream census use.) This premise is stated as the FH-1 restriction of the SC-2 world-to-framework occupancy map, and it is recorded as a **new FH-1 occupancy premise requiring future discharge**: an occupancy map locating the known channels does not by itself establish that no unoccupied admissible stratum lies below the hinge, and whether SC-2 proves that exclusion is an open verification (debt D14). Until that verification, the unconditional content of the paper is "first occupied gap" (P6 alone); "first admissible gap" holds additionally under P6A.

Falsifier (world-facing). Demonstration of an admissible substrate stratum strictly between the floor and the first occupied support — for example, a census extension exhibiting an allowed χ value with no up-type occupant. In that case charm witnesses a later increment, not the first, and the central identity of Theorem 6 fails while every other premise holds.

2.7 P7 — Positive χ-magnitude premise

Where a positive χ metric, spectral distance, or access functional is supplied — the calibration fixing of D2 — it assigns to every pair of distinct admissible strata a positive invariant scalar magnitude:

$\delta\_\chi = \|\Delta\_\chi(U, U')\|\_\chi > 0$.

P7 does not assert that the calibration is fixed; it types what a fixed calibration delivers. The existence claim is conditional exactly as Theorem 2 makes it: without the fixing, only ordinal rank survives (F43).

Falsifier. A magnitude is claimed for the first gap but only an ordinal rank is supplied.

2.8 P8 — Oriented mass-bridge monotonicity premise

Access magnitude and obstruction cost are distinct, oppositely oriented quantities and may not be used interchangeably. With δ_χ the access magnitude and $\mathcal{A}_\chi$ the obstruction cost, the access-to-mass functional satisfies

$$d\mathcal{F}_\chi / d\delta_\chi > 0 \text{ and } d\mathcal{F}_\chi / d\mathcal{A}_\chi < 0 ,$$

and increasing access magnitude cannot reduce admissible mass support unless a separately named suppression factor is present.

Falsifier. The two orientations are conflated, or a higher-access channel is assigned lower admissible mass support without a named suppression bridge.

2.8′ P8A — Net mass-order preservation premise

P8 orders only the χ factor. The full channel coupling factorises (P9) as

$$y_i(\mu) = K_i(\mu) \cdot f_\chi(\lambda_i) , K_i(\mu) = Z_{L,i}^{(-1/2)}(\mu) \cdot B_{u,i} \cdot Z_{R,i}^{(-1/2)}(\mu) ,$$

and the channel-dependent non-χ factors K_i could in principle overturn the f_χ ordering. P8A states that they do not: for each adjacent pair $i = 0, 1$,

$$| \log(f_\chi(\lambda_{i+1}) / f_\chi(\lambda_i)) | > | \log(K_i(\mu) / K_{i+1}(\mu)) | ,$$

equivalently $K_{i+1} f_\chi(\lambda_{i+1}) > K_i f_\chi(\lambda_i)$. Without P8A, χ order guarantees only χ order, not physical mass order, and Stage 2 of P13 has no licence (Theorem 6A).

Falsifier (world-facing). A channel-dependent normalisation, carrier, or dressing log-ratio exceeding the corresponding log-f_χ gap — an order reversal $y_1 < y_0$ or $y_1 > y_2$ traceable to the non-χ factors (F45).

2.9 P9 — Separability premise (operator-level factorisation)

At FH-1 order the effective up-type Yukawa operator is block-diagonal in the physical channel decomposition and multiplicatively renormalised:

$$Y_u(\mu) = Z_L^{(-1/2)}(\mu) \cdot B_u \cdot f_\chi(\hat{\chi}) \cdot Z_R^{(-1/2)}(\mu) ,$$

so that in a χ eigenchannel U_i the scalar coupling factorises as

$$y_i(\mu) = Z\_L,_i\text{^}(-1/2)(\mu) \cdot B_{u,i} \cdot f\_\chi(\lambda_i) \cdot Z\_R,_i\text{^}(-1/2)(\mu) .$$

The scalar factorisation follows from the operator form only if the non-χ factors preserve the χ eigenspaces. P9 therefore requires, at FH-1 order, that B_u, Z_L, and Z_R commute with every spectral projector P_i of $\hat{\chi}$:

$$[P_i , B_u] = [P_i , Z\_L] = [P_i , Z\_R] = 0 ,$$

so that $P_i \cdot Y_u(\mu) \cdot P_i = Z\_L,_i\text{^}(-1/2) \cdot B_{u,i} \cdot f\_\chi(\lambda_i) \cdot Z\_R,_i\text{^}(-1/2) \cdot P_i$ is genuinely derived, with no off-diagonal interference terms. "Block-diagonal in the physical channel decomposition" means block-diagonal in the χ spectral decomposition, with one-dimensional physical blocks.

Corrections are typed. Multiplicative diagonal corrections may enter the dressing factor $\mathcal{D}_i$. Nonfactorisable, channel-mixing corrections may not be absorbed into any scalar factor; they are carried as an explicit operator remainder

$$Y_u(\mu) = Z\_L\text{^}(-1/2)(\mu) \cdot B_u \cdot f\_\chi(\hat{\chi}) \cdot Z\_R\text{^}(-1/2)(\mu) + \Delta Y^{nf}(\mu) , \| \Delta Y^{nf}(\mu) \| < \varepsilon\_nf ,$$

with ε_nf pre-declared and smaller than the relevant singular-value and mass-order gaps, since an operator-valued remainder can shift eigenvectors, disturb the ordering, or violate the alignment premise.

Falsifier. The four-factor decomposition is asserted while the projector-commutation conditions fail; renormalisation is not multiplicative; a nonfactorisable channel-mixing correction is absorbed into a scalar dressing factor; or $\|\Delta Y^{nf}\|$ meets or exceeds the declared bound (F44).

2.10 P10 — Three-channel premise with substrate discriminators

The physical up-type channel space $\mathcal{U}$_phys is three-dimensional. The discriminators are functionals, not names:

$$\Phi(U) = \langle U | \hat{F}\_floor | U \rangle , \Sigma(U) = \langle U | \hat{S}\_end | U \rangle , \Gamma(U_i) = \min\{ \lambda_i - \lambda_{i-1} , \lambda_{i+1} - \lambda_i \} / (\lambda_2 - \lambda_0) ,$$

where Γ is fully constructed from the spectrum of $\hat{\chi}$, and F̂_floor and Ŝ_end are substrate operators — floor-projection and endpoint-saturation functionals — that must be derived or inherited independently of the mass spectrum; their construction is logged as debt D15. P10 states separation inequalities, from which uniqueness and exclusivity follow rather than being assumed:

$$\Phi(U\_F) - \max\{ \Phi(U) : U \neq U\_F \} \geq \Delta\_\Phi > 0 , \Sigma(U\_T) - \max\{ \Sigma(U) : U \neq U\_T \} \geq \Delta\_\Sigma > 0 ,$$
$$U\_T \neq U\_F , \Gamma(U_*) - \max\{ \Gamma(U) : U \neq U_* \} \geq \Delta\_\Gamma > 0 ,$$

with U_* the remaining interior channel. The admissibility thresholds Φ_max, Σ_max, Γ_min are then structurally constrained rather than arbitrary: any threshold chosen inside its separation window (for example, mid-gap) yields the same exclusions, so pre-declaration selects a value within a window the structure has already fixed. The exclusivity consequences: exactly one Φ-excluded channel (U_F), exactly one distinct Σ-excluded channel (U_T), no channel excluded by both, and $\Gamma(U_*) \geq \Gamma\_min$. These fix the standing notation U_F, U_T, and (via §3.3) U_*. Until $\hat{F}\_floor$ and $\hat{S}\_end$ are constructed, Theorem 6 is a conditional discriminator-selection theorem, not a substrate derivation.

Falsifier (audit). Roles or thresholds assigned or tuned after inspection of the mass spectrum; or the selection is described as a substrate derivation while $\hat{F}\_floor$ and $\hat{S}\_end$ remain unconstructed (D15 open).

Falsifier (world-facing). A compressed or degenerate spectrum violating a separation inequality — one channel maximising both Φ and Σ, or the interior channel failing $\Gamma \geq \Gamma\_min$ — emptying the feasible set of the selection rule.

2.11 P11 — Floor-discriminator premise

Exactly one channel satisfies $\Phi(U) \geq \Phi\_max$. A floor-proximate channel is entangled with light-sector suppression, residual access, and confinement-sensitive normalisation, and cannot serve as a clean first-gap magnitude witness without additional light-sector bridges.

Falsifier (audit). No channel is floor-proximate under Φ, or the floor channel is used as a clean $\delta\_\chi{,}_1$ calibrator while light-sector debts remain open.

Falsifier (world-facing). A demonstration that light-quark mass support is scheme-stable and confinement-insensitive — floor entanglement absent — so that the floor channel could after all serve as a clean magnitude witness.

2.12 P12 — Endpoint-discriminator premise

Exactly one channel satisfies $\Sigma(U) \geq \Sigma\_max$. A saturation channel tests terminal access or heavy-sector closure, not first lift-off from the floor.

Falsifier (audit). No channel is saturation-dominated under Σ, or the endpoint channel is used as a first-gap witness while being described as terminal or saturated.

Falsifier (world-facing). A demonstration that the heavy channel's support behaves as an interior step — no saturation signature, e.g. endpoint support responding to access magnitude with the same functional form as the hinge channel.

2.13 P13 — Two-stage identification bridge

The map from internal channels to Standard Model quark labels proceeds in two declared stages.

Stage 1 — sector, by representation content. The bridge datum

$$B\_SM(U_i) = [\ SU(3)\_c\ ,\ SU(2)\_L\ ,\ Y\ ,\ Q\ ,\ \text{chirality}\ ]_i\ ,$$

with the up-type assignment $(3, 2, +1/6)\_L \oplus (3, 1, +2/3)\_R$, identifies each channel as an up-type quark channel. Because the three up-type generations carry the identical representation datum — that triplication is the flavour puzzle itself — Stage 1 fixes the sector and cannot select among the three channels.

Stage 2 — channel, by alignment and order transfer. P17A pairs each χ eigenchannel with a mass eigenchannel; the net mass-order preservation of Theorem 6A — premised in P8A: the non-χ factors do not overturn the f_χ ordering — transfers the χ ordering of the eigenchannels to the mass ordering of their paired eigenchannels; and "charm" is, by Standard Model definition, the intermediate-mass up-type mass eigenchannel. The first-lifted χ channel U_* therefore pairs with the intermediate-mass eigenchannel and is named charm.

Mass ordering thus enters exactly once — at the definitional naming step of Stage 2, as the Standard Model's own naming convention — and never as a selection input. The selection of U_* (§3.3) remains mass-blind. The identification must be unchanged under admissible basis transformations, and it must be declared in which basis U_* is specified — here, as a χ eigenchannel (§3.4), with P17A supplying the alignment to the mass eigenchannel at FH-1 order.

Falsifier. The hinge is named charm without both stages — in particular, if a charm-specific representation datum is claimed (none exists), if mass ordering is used upstream in the selection rather than at the definitional naming step, or if the identified state changes under an admissible basis transformation.

2.14 P14 — Scheme-declared Yukawa premise

The charm-channel target is a scheme-declared, dimensionless Yukawa $y\_C(\mu)$ at declared scale μ, or the equivalent running mass via

$$m\_C(\mu) = (\ v(\mu)\ /\ \sqrt{2}\ ) \cdot y\_C(\mu)\ .$$

Falsifier. An exact charm mass is claimed without scale, scheme, and running conventions.

2.15 P15 — QCD/dressing separation premise

QCD running, confinement dressing, and threshold effects are not χ structure. They are carried by the dimensionless dressing factor $\mathcal{D}\_C(\mu)$.

Falsifier. Running effects are absorbed into $\delta\_\chi,_1$ while $\delta\_\chi,_1$ is claimed as pure access magnitude.

2.16 P16 — Non-insertion premise

The charm Yukawa may not be imported as an independent fitted value. It must be derived through the magnitude-bridge or marked as external calibration. The formal enforcement is the parameter partition of Theorem 10.

Falsifier. $\delta_{\chi,1}$, $\mathscr{F}_\chi$, or any normalisation factor is chosen by fitting y_C, and charm is then declared predicted.

2.17 P17 — CKM-separation premise

Mass eigenvalue support and mixing support are distinct. FH-1 concerns the charm-channel mass magnitude, not CKM angles or phases.

Falsifier. CKM data are used to fix the charm mass channel without a mixing bridge, or CKM closure is claimed from FH-1 alone.

2.18 P17A — χ/mass alignment premise

At FH-1 order, the physical χ-access operator and the Hermitian up-type mass-support operator are simultaneously diagonalisable,

$$[\, \hat{\chi}\, ,\, Y_u^\dagger Y_u \,] = 0\ ,$$

and FH-1 identification statements consume the exact-commutation case.

Approximate alignment ($\varepsilon_{align} > 0$) is admissible only where a declared spectral-projector bound is supplied: a commutator bound alone does not guarantee eigenvector proximity, since the meaning of "small" is fixed by the spectral gaps of both operators. With $M_u = Y_u^\dagger Y_u$ and P_i^χ, P_i^M the relevant spectral projectors, the required bound has the Davis–Kahan form [3]

$$\| P_i^\chi - P_i^M \| \lesssim \| [\, \hat{\chi}\, ,\, M_u \,] \| \,/\, (\, g_\chi \cdot g_M \,)\ ,$$

up to the hypotheses and constants of the invoked perturbation theorem, where g_χ and g_M are the relevant spectral separations. The derivation and application of that bound is part of debt D13; within FH-1, exact commutation corresponds to no mixing, and the controlled failure of commutation is the seed of the downstream CKM gate (D13).

Falsifier. A basis rotation can move the first χ gap from the charm mass eigenchannel to another linear combination, so that the charm-hinge identification is not invariant.

2.19 P18 — No-retrodictive-labelling premise

Channel roles, discriminator thresholds, and identifications may not be assigned or tuned after examining observed masses.

Falsifier. The hinge is identified as charm solely because charm has the observed intermediate mass, or Φ_max/Σ_max are adjusted to make the selection come out right.

2.20 P19 — Rare-decay/mass-hierarchy separation premise

Charm mass-hierarchy support and charm-loop rare-decay support are distinct objects. The hinge channel concerns the first χ gap in the up-type mass ladder; nonlocal charm-loop contributions in $b \to s\, \mu^+\mu^-$ decays concern hadronic loop amplitudes and must be audited through a separate rare-decay transition bridge.

Falsifier. Rare B-decay charm-loop effects are used as evidence for the charm-hinge theorem, or the charm-hinge theorem is used to claim control over $b \to s\, \mu^+\mu^-$ observables without an effective-operator bridge.

2′. Premise-Dependency Map

Result	Depends on premises
Theorem 1 — χ-Gap Necessity	P2, P3, P5, P6
Theorem 2 — First-Gap Magnitude	P2, P3, P4, P6, P7
Theorem 3 — Mass-Bridge Factorisation	P1, P4, P8, P9, P14, P15
Theorem 4 — Light-Floor Exclusion	P10, P11, P14, P15
Theorem 5 — Endpoint Exclusion	P10, P12
Theorem 6 — Hinge Selection and Identification	P2–P13 (incl. P6A, P8A via Theorem 6A), P17A, P18
Theorem 6A — Net Mass-Order Preservation	P8, P8A, P9
Theorem 7 — Charm Non-Insertion	P1–P9, P14–P16; Theorem 6 chain
Theorem 8 — Rare-Decay Separation	P4, P17, P19, I4, I10
Theorem 9 — Hinge Stability	P6, P10, P17A
Theorem 10 — No-Parameter-Hiding	P9, P16, P18

The load-bearing premises are P6 (well-founded spectrum), P6A (lowest-strata occupancy), P8A (net mass-order preservation), P9 (separability with projector commutation), P10–P12 (discriminators and separation inequalities; construction open, D15), P13 (two-stage bridge), and P17A (exact alignment). A referee attacking FH-1 should attack one of these directly.

---

3. Formal Setup: Spectral χ, Discriminators, Closure Levels, and Operator Bridge

3.1 Channel space and spectral χ

Let $\mathcal{U}$ be the up-type flavour-channel space after charge, weak-block, and physical-sector projections, with physical subspace $\mathcal{U}$_phys three-dimensional (P10). By P6, χ-access is carried by the positive self-adjoint operator

$\hat{\chi}$_phys = P_phys · $\hat{\chi}$ · P_phys ,

with well-founded physical spectrum

Spec_phys($\hat{\chi}$) ∋ $\lambda_0 < \lambda_1 < \lambda_2$,

where λ_0 is the floor value and $\lambda_1 = \min S_+(\lambda_0)$. The symbol χ(U) denotes the spectral value of the stratum containing U. χ is not mass; the relation to mass is bridged in §3.5.

Two support sets must be distinguished. $\mathcal{X}$ is the χ-support value space carrying the order ⪯_χ; its physical-sector subset

$\mathcal{X}$_phys ⊆ $\mathcal{X}$

is the set of substrate-admissible support values surviving the positivity and physical-sector projections of I1 and I7. The realised spectrum is the **occupied** subset,

Spec_phys($\hat{\chi}$) = χ($\mathcal{U}$_phys) ⊆ $\mathcal{X}$_phys ,

and the two need not coincide: the substrate may admit strata that no physical channel occupies. The relation between the admissible and occupied sets is fixed by the occupancy premise P6A and by Lemma 1 (§3.3), and is exactly the point at which a silent assumption would otherwise enter.

3.2 Displacement and invariant magnitude

For channels U, U′ define the typed displacement Δ_χ(U, U′) — the ordered movement between their χ strata — and its invariant scalar magnitude

δ_χ = ‖Δ_χ(U, U′)‖_χ = d_χ(χ(U), χ(U′)) .

The first-gap objects are

Δ_χ(U_F, U∗) and δ_χ,₁ = ‖Δ_χ(U_F, U∗)‖_χ = $\lambda_1 - \lambda_0$, the last equality holding in the spectral representation. Coordinate subtraction of χ values is never assumed; where χ is not globally affine, $\delta\_\chi{,}_1$ is the spectral or geodesic distance between the lowest two admissible strata. The legacy notation $\Delta\chi_1$ is reserved for a tangent or coordinate representative of $\Delta\_\chi(U\_F, U*)$ where an affine chart exists.

3.3 Substrate discriminators and the selection rule

By P10, three invariant substrate-level functionals are defined on $\mathcal{U}$_phys with no reference to observed masses or generation names — Γ constructed from the spectrum, Φ and Σ as expectation values of substrate operators whose construction is debt D15:

Φ(U) — floor proximity; Γ(U) — interior isolation; Σ(U) — endpoint saturation;

with thresholds Φ_max, Σ_max, Γ_min pre-declared within the separation windows that P10's inequalities structurally fix. The first clean hinge is defined by the selection rule

$U* = \text{argmin} \{ d\_\chi(\chi(U), \chi_0) > 0 : \Phi(U) < \Phi\_\text{max} , \Sigma(U) < \Sigma\_\text{max} , \Gamma(U) \geq \Gamma\_\text{min} \}$,

the argmin taken over U ∈ $\mathcal{U}$_phys: among channels strictly lifted above the floor that are neither floor-proximate nor saturation-dominated and are interior-isolated, U∗ is the one at least lifted distance. The Γ ≥ Γ_min condition is the cleanliness clause: it certifies strictly positive adjacent gaps around the hinge stratum, which is what makes the stability bound of Theorem 9 nonvacuous. The selection consumes only ($\hat{\chi}$, Φ, Γ, Σ) and the pre-declared thresholds. No mass datum enters.

Only after U∗ is fixed does the two-stage bridge of P13 identify it with a Standard Model label. This is the structural difference between

"the middle channel is charm, so charm is the hinge"

and

"VERSF selects an interior, first-lifted channel without using mass data; the two-stage bridge (P13) identifies that selected state as charm."

Lemma 1 — Occupancy identity. Under P6 and P6A,

$\chi(U*) = \lambda_1 = \min \{ \chi \in \mathcal{X}\_\text{phys} : \chi_0 \prec\_\chi\, \chi \}$.

Proof. By P6A the occupied spectrum consists of the three $\preceq\_\chi$-least elements of $\mathcal{X}$_phys, so no admissible stratum lies strictly between $\lambda_0 = \chi_0$ and λ_1. The minimum of the admissible upper set above the floor therefore coincides with the minimum of the occupied upper set, which by P6 is λ_1; and by the selection rule $\chi(U*) = \lambda_1$. ■

Without P6A only the weaker identity holds — U∗ witnesses the first **occupied** increment — and the symbol λ_1 would silently carry two inequivalent readings (spectral minimum and admissible minimum). Lemma 1 is what licenses the single-symbol usage throughout the paper.

3.4 Basis declaration and alignment

Before diagonalising the Yukawa operator, "charm" is basis-dependent: a weak-interaction channel, a χ eigenchannel, and a mass eigenchannel are not automatically identical. FH-1 declares that U∗ is specified as a **χ eigenchannel**, and P17A supplies the alignment between χ eigenchannels and mass eigenchannels at FH-1 order — exactly ($[\hat{\chi}, Y_u^\dagger Y_u] = 0$). FH-1 identification statements consume the exact case only; approximate alignment is admissible solely under the declared Davis–Kahan spectral-projector bound of P17A and is deferred, with the controlled failure of alignment, to the CKM gate as debt D13.

3.5 The dimensionless factorised bridge

An admissible charm-channel Yukawa claim has the form

$$y\_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathscr{R}\_H(\mu) \cdot \mathscr{F}\_\chi(\delta\_\chi{,}_1) \cdot \mathcal{D}\_C(\mu) ,$$

with all four factors declared dimensionless normalised quantities, so that $y\_C(\mu)$ is dimensionless as a Yukawa must be:

$$[\mathcal{N}_u] = [\mathscr{R}\_H] = [\mathscr{F}\_\chi] = [\mathcal{D}\_C] = 1 .$$

In particular $\mathscr{R}\_H$ is not a field amplitude, vacuum expectation value, or stiffness; it is the normalised carrier-coupling ratio

$$\mathscr{R}\_H(\mu) = g\_H{,}U(\mu) / g\_H{,}ref(\mu) .$$

The dimensionful mass statement, when required, is carried entirely by the standard relation

$$m\_C(\mu) = (v(\mu) / \sqrt{2}) \cdot y\_C(\mu) .$$

Reduced support is

$$\hat{y}\_C(\mu) = y\_C(\mu) / [\mathcal{N}_u(\mu) \cdot \mathscr{R}\_H(\mu) \cdot \mathcal{D}\_C(\mu)] = \mathscr{F}\_\chi(\delta\_\chi{,}_1) ,$$

and where $\mathscr{F}\_\chi$ is positive the log-access form $\Lambda\_\chi{,}_1 = \log \mathscr{F}\_\chi(\delta\_\chi{,}_1)$ expresses the first gap as an additive access magnitude.

3.6 Operator-level grounding of the bridge

The four scalar factors are not postulated. By P9, the effective action contains

$\mathscr{L}$_eff ⊃ − $\bar{Q}$_L · $Y_u(\mu)$ · $\tilde{H}$ · U_R + h.c. ,

with block-diagonal, multiplicatively renormalised effective Yukawa operator

$Y_u(\mu)$ = Z_L^(−1/2)(μ) · B_u · f_χ(χ̂) · Z_R^(−1/2)(μ) .

Projection onto the χ eigenchannel U_i (legitimate at FH-1 order by P17A) yields

$y_i(\mu)$ = Z_L,$_i$^(−1/2)(μ) · $B_{u,i}$ · f_χ(λ_i) · Z_R,$_i$^(−1/2)(μ) ,

and the identifications $\mathcal{N}_u$ · $\mathscr{R}$_H ↔ Z-factors and carrier normalisation, $\mathscr{F}$_χ(δ_χ,$_1$) ↔ f_χ(λ_1) relative to the floor reference, and $\mathcal{D}$_C ↔ the declared running/threshold corrections, give the scalar factorisation as a projection of operator structure rather than an ansatz.

3.7 Orientation of magnitude and cost

Access magnitude δ_χ and obstruction cost $\mathcal{A}$_χ carry opposite monotonic orientations (P8):

d$\mathscr{F}$_χ / dδ_χ > 0 , d$\mathscr{F}$_χ / d$\mathcal{A}$_χ < 0 .

Representations such as $\mathscr{F}$_χ(δ_χ) = e^(−$\mathcal{A}$_χ(δ_χ)) are admissible provided the cost $\mathcal{A}$_χ is a decreasing function of the magnitude δ_χ, so that the composite remains increasing in δ_χ. The two words are never used interchangeably in this paper.

3.8 Three distinct closure levels

To prevent ordinal ordering, metric separation, and numerical mass prediction from being conflated, FH-1 distinguishes three logically different closure levels.

Level I — Ordinal χ closure. An invariant ordering exists:

U_F ≺_χ U∗ ≺_χ U_T .

This determines which channel lies above which, but not the physical separation between them.

Level II — Metric χ closure. The ordered support carries an independently defined invariant distance or access cost:

d_χ(χ(U_F), χ(U∗)) = δ_χ,$_1$ > 0 .

This establishes that the first lifted channel is physically separated from the floor by a nonzero, convention-independent magnitude. It does not yet determine a Yukawa coupling.

Level III — Dynamical mass-magnitude closure. A substrate-derived map converts the invariant separation into reduced mass support:

$y\_C(\mu) / y\_ref(\mu) = \mathcal{M}\_\chi(\delta\_\chi,_1 ; \lambda\_substrate)$,

where the map $\mathcal{M}\_\chi$ and every parameter in λ_substrate are fixed independently of the observed charm mass.

FH-1 conditionally closes Level I, specifies exactly what Level II requires, and constructs the audit architecture of Level III. It does not claim numerical Level III closure unless the corresponding bridge debts (D1–D7) have been discharged. The word "Magnitude-Bridge" in the title refers to this architecture: the paper establishes the necessary carrier of magnitude, not yet its value.

---

4. Why a Mass Carrier Does Not Derive a Hierarchy

The Higgs-radial gate supplies a scalar mass carrier: it explains how mass support can exist as closure-interface stiffness. But a carrier is not a hierarchy. A single carrier can support many channels; it does not determine why one channel receives tiny support, another intermediate, another large.

The Standard Model permits writing independent Yukawa parameters y_u, y_c, y_t. That is valid phenomenology, not derivation. A VERSF derivation must not write three independent values and then say "χ explains them." To explain them, χ must be defined before the values are known, and the bridge from χ to mass support must be stated. Charm is the first serious test of that requirement: the floor channel is too entangled with light-sector effects to expose first-gap structure, and the endpoint channel is governed by saturation. The interior channel is where the first positive separation is visible.

---

5. From Ordering to a First Gap

By P3 and P5, the physical spectrum contains at least one strict inequality. By P6, the positive upper set above the floor,

$S_+(\lambda_0) = \{ \lambda \in Spec\_phys(\chi) : \lambda > \lambda_0 \}$,

is nonempty and possesses a least element λ_1. This is stronger than bare discreteness: a discrete set need not have a least point above λ_0 unless the relevant upper set has a minimum, and P6 requires exactly that. The first gap is then the spectral object

$\delta\_\chi,_1 = \lambda_1 - \lambda_0 > 0$, not a verbal ladder step. P6A extends this from the occupied spectrum to the admissible support set: by Lemma 1, the first occupied gap is also the first admissible gap. If the first gap were not named, the hinge channel could be treated as arbitrary; if named without magnitude, it would remain ordinal; if its magnitude were fitted from charm, the derivation would be circular. FH-1 therefore requires all three: ordered spectrum, well-founded first gap, and independent magnitude — the content of Theorems 1 and 2.

---

6. From Gap to Invariant Magnitude

A rank ordering u < c < t does not say how far the hinge is above the floor. The magnitude is supplied by the invariant distance

$$\delta_{\chi,1} = \|\Delta_\chi(U_F, U_*)\|_\chi = d_\chi(\chi_0, \chi_1) ,$$

which must be positive (else the hinge is not separated from the floor), invariant (else it is a label artifact), and defined on χ structure rather than observed mass (else it is a fit). The access-to-mass functional $\mathscr{F}_\chi$ then carries the oriented monotonicity of P8: increasing in access magnitude, decreasing in obstruction cost. The central non-insertion rule is that $\mathscr{F}_\chi$ and $\delta_{\chi,1}$ are fixed before any comparison with the charm-channel Yukawa; if either is chosen to reproduce charm, FH-1 is a reparametrisation, not a prediction. Theorem 10 formalises this as a parameter-partition and identifiability condition.

---

7. The Mass-Bridge Factorisation

A fermion mass-channel value contains at least four distinct layers: carrier support, χ magnitude, sector normalisation, and scheme/running/dressing. If these are not separated, any one can hide a fit. The factorised bridge

$$y_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathscr{R}_H(\mu) \cdot \mathscr{F}_\chi(\delta_{\chi,1}) \cdot \mathcal{D}_C(\mu)$$

separates them, and §3.6 grounds the separation at operator level so it is a projection of block-diagonal structure, not an ansatz.

Division of labour: $\mathscr{R}_H$ supplies carrier support and may not encode hierarchy; $\mathscr{F}_\chi$ supplies first-gap magnitude and may not absorb running or Higgs normalisation; $\mathcal{N}_u$ supplies sector-level normalisation and may not be chosen per channel; $\mathcal{D}_C$ supplies running, threshold, and dressing corrections and may not become a hidden fitted Yukawa.

The floor channel's reduced support $\hat{y}_F$ may contain residual access, chiral suppression, confinement-sensitive normalisation, and boundary effects; the endpoint channel's $\hat{y}_T$ is governed by saturation and terminal constraints. Neither isolates $\mathscr{F}_\chi(\delta_{\chi,1})$. The interior channel's reduced support is the clean target:

$$\hat{y}_C(\mu) = \mathscr{F}_\chi(\delta_{\chi,1}) .$$

---

8. Floor and Endpoint Exclusions by Discriminator

8.1 The floor problem

The channel with $\Phi(U) \geq \Phi_{max}$ has its χ access compressed against the light-sector boundary. Its apparent mass support may be determined by residual closure leakage, chiral or near-chiral suppression, confinement-sensitive normalisation, weak attachment thresholds, and sector-floor regularisation — effects not identical to the first χ gap. It is therefore not a clean witness of $\delta_{\chi,1}$.

8.2 The endpoint problem

The channel with $\Sigma(U) \geq \Sigma_{max}$ is governed by maximum access, endpoint saturation, heavy-sector closure, and proximity to electroweak-scale support. It measures how high the ladder goes, not the first step off the floor.

8.3 The hinge criterion

A clean first-gap witness must satisfy four checkable conditions: it is strictly lifted above the floor; it lies below the Φ and Σ thresholds and at or above the Γ_{min} isolation threshold; its χ value equals $\min\{ \chi \in \mathcal{X}_{phys} : \chi_0 \prec_\chi \chi \}$ — checkable via Lemma 1 under P6A, rather than assumed; and it carries reduced support as $\mathscr{F}_\chi(\delta_{\chi,1})$. In the three-channel space, the selection rule of §3.3 shows exactly one channel qualifies.

---

9. Hinge Selection and the Charm Bridge

9.1 Selection before naming

The selection rule

$$U_* = \operatorname{argmin} \{ d_\chi(\chi(U), \chi_0) > 0 : \Phi(U) < \Phi_{max} , \Sigma(U) < \Sigma_{max} , \Gamma(U) \geq \Gamma_{min} \}$$

fixes the hinge from ($\hat{\chi}$, Φ, Γ, Σ) alone. By the selection rule $\chi(U_*) = \lambda_1$; by Lemma 1 under the occupancy premise P6A, λ_1 is also the first admissible support above the floor — not merely the first occupied one.

9.2 Naming after selection

By P13 identification proceeds in two stages. Stage 1: the representation datum

B_SM(U∗) = [SU(3)_c , SU(2)_L , Y , Q , chirality]

with the up-type assignment (3, 2, +1/6)_L ⊕ (3, 1, +2/3)_R fixes U∗ as an up-type quark channel — and can do no more, since all three up-type generations carry this identical datum. Stage 2: P17A pairs U∗ with a mass eigenchannel, Theorem 6A (under P8A) transfers the χ ordering to the mass ordering of the paired eigenchannels, and "charm" is by Standard Model definition the intermediate-mass up-type eigenchannel; hence U∗ ↦ c. Mass ordering enters only here, as definitional naming, and does not retro-contaminate the mass-blind selection of §9.1 — U∗ is not identified as charm because it is "second", nor by any charm-specific representation datum, which does not exist.

9.3 What would count as prediction

FH-1 becomes numerically predictive only if later or prior gates supply: $\delta_{\chi,1}$ from χ substrate structure (D2, D3); $\mathscr{F}_\chi$ from the access-cost law (D4); $\mathscr{R}_H$ from Higgs-radial/electroweak normalisation (D5); $\mathcal{N}_u$ from the up-sector bridge (D6); and $\mathcal{D}_C$ from scheme and running (D7). Then $y_C(\mu)$ follows from the bridge. Without those ingredients, FH-1 closes structural admissibility, not numerical prediction.

---

10. Scheme, Running, Confinement, and Loop Firewalls

A bare charm mass is not the target. The theorem target is the scheme-declared, dimensionless $y_C(\mu)$, or the running mass via $m_C(\mu) = (v(\mu)/\sqrt{2})\cdot y_C(\mu)$.

QCD is not χ. Running effects belong in $\mathcal{D}_C(\mu)$; hiding them inside $\delta_{\chi,1}$ destroys the access meaning of χ.

Confinement is not χ. The floor channel is especially confinement-sensitive — one reason it is excluded as the clean witness. Charm is better, but still a quark: any numerical comparison requires the running/dressing bridge.

Mixing is not χ magnitude. CKM structure concerns the relation between weak and mass eigenstates. FH-1 concerns the hinge-channel magnitude. The controlled failure of P17A alignment is the designed entry point of the later CKM gate (D13).

Loop contamination is not χ magnitude. Charm-loop content in rare amplitudes is a transition-class object and is not a component of any bridge factor (§10″, Theorem 8).

10′. Open Bridge-Debt Ledger

FH-1 licenses downstream flavour work only against the following explicitly recorded debts.

Debt	Content	To be discharged by
D1 — χ-origin debt	Derivation of the χ operator from substrate, commitment, or weak-block primitives	χ-readout foundation gate
D2 — χ-metric debt	Derivation of the positive χ metric or access functional	χ-magnitude gate
D3 — $\delta\_\chi{,}_1$ numerical debt	Exact value of the first physical χ gap	χ-quantisation gate
D4 — $\mathscr{F}\_\chi$ / $\boldsymbol{\mathcal{M}}\_\chi$ law debt	Derivation of the access-to-mass map and its substrate parameters	flavour-access dynamics gate
D5 — Higgs normalisation debt	Numerical carrier normalisation $\mathscr{R}$_H	scalar/electroweak normalisation gate
D6 — up-sector normalisation debt	Derivation of $\mathcal{N}_u(\mu)$	up-sector bridge gate
D7 — QCD/running debt	Derivation of $\mathcal{D}$_C(μ) in a declared scheme	QCD/running bridge
D8 — light-floor debt	Derivation of floor-channel support and suppression factors	light-sector floor theorem
D9 — endpoint debt	Derivation of endpoint saturation support	top-endpoint theorem
D10 — CKM debt	Derivation of mixing angles/phases	mixing-interface theorem
D11 — empirical-comparison debt	Comparison with observed charm data after all conventions fixed	programme audit
D12 — rare-decay transition debt	Loop-level $b \to s\, \mu^+\mu^-$ effective-operator structure, charm-loop hadronic terms, muon-current access, possible Wilson-coefficient shifts	rare-decay / FCNC anomaly gate (FCNC-1)
D13 — χ/mass alignment debt	Derivation of the commutator structure $[\hat{\chi}, Y_u{\dagger}Y_u]$ and the bound ε_align; exact commutation ⇒ no mixing, controlled failure ⇒ mixing	CKM/mixing gate

Debt	Content	To be discharged by
D14 — occupancy debt	Verification that SC-2 establishes the below-hinge exclusion (no unoccupied admissible stratum under the hinge), and derivation of lowest-strata occupancy from the census machinery	SC-2 bridge / census gate
D15 — discriminator-construction debt	Derivation or inheritance of the substrate operators F̂_floor and Ŝ_end defining Φ and Σ, independent of the mass spectrum; Γ is already spectrally constructed	substrate-discriminator gate

No downstream paper may treat D1–D15 as discharged by FH-1.

10″. Rare-Decay Anomaly and Charm-Loop Firewall

FH-1 concerns charm as a mass-hierarchy hinge. It does not concern nonlocal charm-loop effects in rare $b \to s\,\mu^+\mu^-$ transitions, and the word "charm" must not be allowed to blur the two:

- **FH-1 charm** — the hinge (U_*, $\delta\_\chi_{,1}$), a mass-eigenchannel object;
- **B-anomaly charm** — charm-loop contributions, a major theoretical uncertainty in interpreting selected $b \to s\,\mu^+\mu^-$ observables.

Measurements of $B^0 \to K^0$ *$\mu^+\mu^-$ angular observables and branching fractions (LHCb, full 8.4 fb^{-1} dataset including S-wave and muon-mass effects [8]; CMS, 140 fb^{-1} at 13 TeV [9]) report tensions with Standard Model predictions in CP-averaged observables and branching fractions, while CP-asymmetry observables show no significant deviation. A distinct LHCb amplitude analysis on the same luminosity [12] isolates the local and nonlocal contributions across the full dimuon spectrum without charmonium vetoes, finding nonlocal interference larger than predicted but with only minor impact on the fitted Wilson coefficients — the state-of-the-art record on precisely the charm-loop object that debt D12 must bridge. Global analyses frame the pattern as potentially consistent with a shift in the semileptonic Wilson coefficient C_9^μ (operator conventions of [11]), with charm-loop hadronic contributions the principal interpretive uncertainty. The lepton-universality ratios R_K and R_K* agree with the Standard Model [10]; the tension recorded here is therefore concentrated in branching fractions and angular observables rather than in lepton-universality ratios. None of this is FH-1 content; it is the stress-test environment for the later FCNC gate.

The firewall:

Charm as U_*, the first χ-gap mass hinge, may not be identified with charm-loop contamination in $b \to s\,\mu^+\mu^-$ amplitudes.

Any future bridge (debt D12) must distinguish five typed objects — mass-hierarchy charm ($\delta\_\chi,_1$); virtual charm-loop support; b → s flavour-changing neutral-current access; muon-current access; and Wilson-coefficient shifts such as ΔC_9^μ acting on

$\mathcal{O}_9$^μ = (s̄ γ_α P_L b)(μ̄ γ^α μ) , $\mathcal{O}_{10}$^μ = (s̄ γ_α P_L b)(μ̄ γ^α γ_5 μ) —

and, by I10, must treat CP-averaged and CP-asymmetry observables as separate diagnostics, since the observed pattern is asymmetric between them. FH-1 names these objects solely so that D12 has a precise target; it derives none of them.

---

11. Theorem 1 — χ-Gap Necessity Theorem

11.1 Statement

Theorem 1 — χ-Gap Necessity Theorem. If the physical up-type channel space carries a positive, invariant, self-adjoint χ-access operator whose channels are not all χ-degenerate, and whose physical spectrum is well founded in the sense of P6, then a first positive χ gap exists:

$\delta\_\chi,_1 = \lambda_1 - \lambda_0 > 0$.

Depends on: P2, P3, P5, P6.

11.2 Proof

By P2 and P6, χ-access is carried by $\hat{\chi}$_phys = P_phys·$\hat{\chi}$·P_phys with invariant physical spectrum. By P3 the spectral values are invariantly ordered. By P5 not all channels are degenerate, so the spectrum contains at least two distinct values; in particular the positive upper set

$S_+(\lambda_0) = \{ \lambda \in \mathrm{Spec\_phys}(\hat{\chi}) : \lambda > \lambda_0 \}$

is nonempty. By P6 this set possesses a least element λ_1 — the well-foundedness condition, which is strictly stronger than discreteness, since a discrete set need not have a minimum above λ_0 unless its upper set does. Then

$\delta\_\chi,_1 = \lambda_1 - \lambda_0 > 0$

is the first physical χ gap, invariant because the spectrum is invariant. ■

11.3 Corollary — No first gap without well-foundedness

If the positive spectrum accumulates at the floor, no least positive separation exists and FH-1 cannot claim a first gap without a regulator, threshold, or quantisation theorem (falsifier F6).

---

12. Theorem 2 — First-Gap Magnitude Theorem

12.1 Statement

Theorem 2 — First-Gap Magnitude Theorem. The first χ gap carries a convention-independent physical magnitude **if and only if** the spectral calibration of the χ operator — or an equivalent invariant distance — is physically fixed. When it is, the magnitude is

$$\delta_{\chi,1} = \|\Delta_\chi(U_F, U_*)\|_\chi = d_\chi(\chi_0, \chi_1) > 0 \ ,$$

positive, invariant, and — by the χ-independence premise — defined independently of the observed charm mass. Where χ is not globally affine, $\delta_{\chi,1}$ is the spectral or geodesic distance between the lowest two admissible strata, not coordinate subtraction.

Depends on: P2, P3, P4, P6, P7.

12.2 Proof

Sufficiency. By Theorem 1 the first gap exists as an ordered spectral separation. Order alone determines which stratum comes first, not the size of the step. With the spectral calibration fixed — $\hat{\chi}$ determined up to unitary equivalence only (P6), the monotone-recalibration freedom excluded — the eigenvalue difference $\lambda_1 - \lambda_0$ is itself a canonical scalar; no additional metric structure is required beyond the fixing. By P7 the typed displacement $\Delta_\chi(U_F, U_*)$ carries this as the invariant scalar magnitude $\delta_{\chi,1} = \|\Delta_\chi(U_F, U_*)\|_\chi > 0$. Because the measure is invariant, the magnitude is not a label artifact; because it is positive, the hinge stratum is physically separated from the floor; and by P4 the magnitude is defined on χ structure rather than observed mass, so it is not a charm fit. In the spectral representation of P6 the magnitude coincides with $\lambda_1 - \lambda_0$; in non-affine representations it is the declared invariant distance, so no global linear structure on χ is assumed. Invariance here means invariance within the unitary-equivalence class declared by P6: a monotone spectral recalibration $\hat{\chi} \to f(\hat{\chi})$ would change the gap while preserving every ordinal fact, and the exclusion of that freedom — the physical fixing of the χ scale — is routed to debt D2, on which Level II invariance is conditional (F43).

Necessity. Suppose the spectral calibration is not physically fixed. Then the monotone recalibration $\hat{\chi} \to f(\hat{\chi})$ is a permitted relabelling: it preserves the ordering, the eigenvectors, and every ordinal fact while changing every gap, so any candidate magnitude assignment changes under permitted transformations and no convention-independent magnitude exists — this is exactly the freedom routed to D2 (F43). In the limiting case where only the order $\preceq_\chi$ survives, two positive increments of equal ordinal rank are indistinguishable under order-preserving relabellings, contradicting the invariance required by P3 and I7: ordinal rank is all that survives. ■

12.3 Corollary — Rank is not magnitude

A generational label such as "second generation" supplies order only. A magnitude claim requires the invariant distance; choosing that distance to reproduce charm renders the theorem circular (F21).

---

13. Theorem 3 — Mass-Bridge Factorisation Theorem

13.1 Statement

Theorem 3 — Mass-Bridge Factorisation Theorem. Under the separability premise, an admissible charm-channel Yukawa arises as the χ-eigenchannel projection of a block-diagonal, multiplicatively renormalised effective Yukawa operator, and therefore factorises as the dimensionless bridge

$$y_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathcal{R}_H(\mu) \cdot \mathcal{F}_\chi(\delta_{\chi,1}) \cdot \mathcal{D}_C(\mu) .$$

A bare charm Yukawa inserted without this structure is not derived.

Depends on: P1, P4, P8, P9, P14, P15.

13.2 Proof

By P9 the effective action contains $\mathcal{L}_{eff} \supset -\bar{Q}_L \cdot Y_u(\mu) \cdot \tilde{H} \cdot U_R$ + h.c. with

$$Y_u(\mu) = Z_L^{(-1/2)}(\mu) \cdot B_u \cdot f_\chi(\hat{\chi}) \cdot Z_R^{(-1/2)}(\mu) ,$$

block-diagonal in the χ spectral decomposition and multiplicatively renormalised, with B_u, Z_L, Z_R commuting with every χ spectral projector P_i (P9). Projecting onto the χ eigenchannel U_i therefore gives, with no off-diagonal interference,

$$y_i(\mu) = Z_{L,i}^{(-1/2)}(\mu) \cdot B_{u,i} \cdot f_\chi(\lambda_i) \cdot Z_{R,i}^{(-1/2)}(\mu) .$$

By P1 the carrier dependence enters through the Higgs-radial factor, isolated in the normalised ratio $\mathcal{R}_H(\mu) = g_{H,U}(\mu)/g_{H,ref}(\mu)$. By P4 the χ dependence enters only through f_χ evaluated on the spectrum; relative to the floor reference this is $\mathscr{F}\_\chi(\delta\_\chi_{,1})$, oriented per P8. The remaining Z-factor and sector content define $\mathcal{N}_u(\mu)$, and by P14–P15 the scheme, running, and threshold content is carried by $\mathcal{D}\_C(\mu)$, all factors dimensionless with y_C dimensionless and $m\_C(\mu) = (v(\mu)/\sqrt{2})\cdot y\_C(\mu)$.

If a paper writes y_C as a bare fitted number, the sources of carrier support, χ magnitude, normalisation, and running are indistinguishable, and nothing has been derived. Where corrections arise beyond FH-1 order, P9 types them: multiplicative diagonal corrections enter $\mathcal{D}\_C$; nonfactorisable channel-mixing corrections remain in the bounded operator remainder ΔY^{nf} ($\|\Delta Y^{nf}\| < \varepsilon\_nf$) and may not be absorbed into any scalar factor. The factorisation is exact at gate order and auditable beyond it. ■

13.3 Corollary — Reduced support is the clean target

$\hat{y}\_C(\mu) = y\_C(\mu)/[\mathcal{N}_u(\mu)\cdot\mathscr{R}\_H(\mu)\cdot\mathcal{D}\_C(\mu)] = \mathscr{F}\_\chi(\delta\_\chi_{,1})$; any factor tuned specifically to reproduce charm functions as a hidden Yukawa unless independently derived (Theorem 10).

---

14. Theorem 4 — Light-Floor Exclusion Theorem

14.1 Statement

Theorem 4 — Light-Floor Exclusion Theorem. The channel with $\Phi(U) \geq \Phi\_max$ cannot serve as the clean witness of the first χ gap: its mass support is entangled with light-sector suppression, confinement-sensitive normalisation, and residual access effects unless additional light-floor bridges are supplied.

Depends on: P10, P11, P14, P15.

14.2 Proof

By P10–P11 exactly one channel is floor-proximate under the pre-declared threshold, and floor proximity is a substrate condition, not a mass label. A floor-proximate channel can receive apparent support from residual floor access, chiral suppression, confinement-sensitive normalisation, threshold and scheme effects, and sector-floor regularisation — none identical to $\delta\_\chi_{,1}$. By P14–P15 its mass value is scheme-dependent and dressing-entangled at exactly the point where those effects are largest. Hence its reduced support cannot isolate $\mathscr{F}\_\chi(\delta\_\chi_{,1})$ without a separate light-floor theorem (debt D8), and it is excluded as the clean witness. ■

14.3 Corollary — Floor effects may not be folded into $\delta\_\chi_{,1}$

Absorbing floor suppression into the gap destroys the separation between χ magnitude and floor physics (F13, F32).

---

15. Theorem 5 — Endpoint Exclusion Theorem

15.1 Statement

Theorem 5 — Endpoint Exclusion Theorem. The channel with $\Sigma(U) \geq \Sigma\_max$ cannot serve as the first-gap witness: it tests saturation, terminal access, or heavy-sector closure, not first lift-off from the floor.

Depends on: P10, P12.

15.2 Proof

By P10 and P12 exactly one channel is saturation-dominated under the pre-declared threshold, and saturation is a substrate condition. Inferring $\delta\_\chi,_1$ from that channel would mix at least two distinct structures — the first gap and endpoint/saturation support — so the endpoint channel cannot isolate the first gap. Its support is deferred to a separate endpoint theorem (debt D9). ■

15.3 Corollary — Saturation is not the first gap

A heavy-sector saturation effect may not be counted as $\delta\_\chi,_1$ (F14).

---

16. Theorem 6 — Hinge Selection and Charm Identification Theorem

16.1 Statement

Theorem 6 — Hinge Selection and Charm Identification Theorem. Under the spectral, discriminator, alignment, and representation-bridge premises, the selection rule

$U_* = \text{argmin} \{ d\_\chi(\chi(U), \chi_0) > 0 : \Phi(U) < \Phi\_max , \Sigma(U) < \Sigma\_max , \Gamma(U) \geq \Gamma\_min \}$

fixes a unique interior first-lifted channel without reference to observed masses or generation names, with $\chi(U_*) = \lambda_1$; and the two-stage bridge of P13 identifies U_* with the charm mass eigenchannel — sector by representation content, channel by P17A alignment and Theorem 6A (P8A) order transfer — with mass ordering entering only as the Standard Model's definitional naming convention, and the identification invariant under admissible basis transformations.

Depends on: P2–P13 (including P6A, P8A via Theorem 6A), P17A, P18.

16.2 Proof

Selection. By Theorem 1 the first gap exists; by Theorem 2 it has invariant magnitude. By P10 the physical space is three-dimensional with pre-declared discriminators and thresholds. By P11 exactly one channel fails the floor condition $\Phi < \Phi\_max$; by P12 exactly one channel fails the saturation condition $\Sigma < \Sigma\_max$; by the separation inequalities of P10 these are distinct channels and no channel fails both, and the remaining interior channel satisfies $\Gamma \geq \Gamma\_min$. The feasible set of the selection rule — channels strictly lifted above the floor satisfying all three conditions — therefore contains exactly one element, which the argmin returns: U_*, with

$$\chi(U_*) = \lambda_1 = \min S_+(\lambda_0) ,$$

the only feasible channel being the least-lifted occupied one. By Lemma 1, under the occupancy premise P6A this occupied minimum coincides with the minimum admissible support above the floor,

$$\chi(U_*) = \min \{ \chi \in \mathcal{X}\_phys : \chi_0 \prec\_\chi\, \chi \} ,$$

so U_* witnesses the first **admissible** increment, not merely the first occupied one. The selection consumed only $(\hat{\chi}, \Phi, \Gamma, \Sigma)$ and the pre-declared thresholds; by P18 none of these may be tuned after inspecting masses, so no mass datum entered.

Identification. By P13 Stage 1, the representation datum B_SM(U_*) — colour, weak isospin, hypercharge, charge, chirality — with the up-type assignment $(3, 2, +1/6)\_L \oplus (3, 1, +2/3)\_R$ fixes U_* as an up-type quark channel. Stage 1 cannot select among the three generations, which carry the identical datum. By P17A, U_* coincides with a mass eigenchannel exactly at FH-1 order (the approximate-alignment case is admissible only under the declared spectral-projector bound of P17A and is deferred to D13); by Theorem 6A under P8A the χ ordering transfers to the mass ordering of the paired eigenchannels — P8 alone orders only the f_χ factor — so the first-lifted χ channel pairs with the intermediate-mass eigenchannel. Since "charm" is by Standard Model definition the intermediate-mass up-type mass eigenchannel, the bridge maps $U_* \mapsto c$. Mass ordering enters only at this definitional naming step — it does not contaminate the selection above, which consumed no mass datum — and "the charm χ value" is unambiguous and basis-invariant.

Therefore VERSF selects an interior, first-lifted channel without using mass data, and the two-stage bridge (P13) identifies that selected state as charm. ■

16.3 Corollary — The hinge is selected, not labelled

The result is not "the middle channel is charm, so charm is the hinge." Charm enters only after U_* is fixed by substrate discriminators; it is the minimal admissible non-floor up-type channel within the three-state χ spectrum.

16′. Theorem 6A — Net Mass-Order Preservation Theorem

16′.1 Statement

Theorem 6A — Net Mass-Order Preservation Theorem. Let $y_i(\mu) = K_i(\mu) \cdot f\_\chi(\lambda_i)$ with $\lambda_0 < \lambda_1 < \lambda_2$, f_χ strictly increasing (P8), and $K_i(\mu)$ the combined non-χ factor of P9. If for each adjacent pair i = 0, 1

$$| \log(f\_\chi(\lambda_{i+1}) / f\_\chi(\lambda_i)) | > | \log(K_i(\mu) / K_{i+1}(\mu)) | \quad \text{(P8A)} ,$$

then $y_0(\mu) < y_1(\mu) < y_2(\mu)$. Consequently the first-lifted χ eigenchannel is paired with the intermediate-mass up-type eigenchannel.

Depends on: P8, P8A, P9.

16′.2 Proof

For each adjacent pair,

$$\log y_{i+1}(\mu) - \log y_i(\mu) = \log(f\_\chi(\lambda_{i+1})/f\_\chi(\lambda_i)) - \log(K_i(\mu)/K_{i+1}(\mu)) .$$

The first term is strictly positive: $\lambda_{i+1} > \lambda_i$ and f_χ is strictly increasing (P8). By P8A the magnitude of the first term exceeds the magnitude of the second, whatever its sign. The difference is therefore strictly positive and $y_{i+1}(\mu) > y_i(\mu)$. Applying this at i = 0 and i = 1 yields the strict chain $y_0(\mu) < y_1(\mu) < y_2(\mu)$. The χ ordering of the eigenchannels thus coincides with the mass ordering of their paired eigenchannels, so the first-lifted χ channel is paired with the intermediate-mass eigenchannel. ■

16′.3 Corollary — the naming licence

Without P8A, Theorem 6 licenses only "U_* is the first χ hinge"; it is Theorem 6A that upgrades this to "U_* is charm." Stage 2 of P13 consumes Theorem 6A, not P8 alone. An order reversal traceable to the non-χ factors is falsifier F45, and it would strip the naming while leaving the selection intact.

---

17. Theorem 7 — Charm Non-Insertion Theorem

17.1 Statement

Theorem 7 — Charm Non-Insertion Theorem. Within the restricted FH-1 audit class, satisfaction of the factorised, independently normalised, non-inserted χ bridge

$$y\_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathcal{R}\_H(\mu) \cdot \mathcal{F}\_\chi(\delta\_\chi{,}_1) \cdot \mathcal{D}\_C(\mu) \,,$$

with $\delta\_\chi{,}_1$ and $\mathcal{F}\_\chi$ fixed independently of the observed charm mass, is **necessary and gate-sufficient** for admitting charm as the first-gap candidate. It is not sufficient for full phenomenological validity of the VERSF flavour sector. Any charm Yukawa imported without this structure is external Standard Model inheritance, calibration, or bookkeeping, not a VERSF derivation.

Depends on: P1–P9, P14–P16, and the Theorem 6 chain (P2–P13 incl. P6A, P17A, P18).

17.2 Proof

Necessity. An admissible charm-channel Yukawa must pass through the carrier (P1), relate to χ structure only through the bridged access functional (P2–P8), factorise per the separability premise (P9, Theorem 3), and be scheme-declared with running separated (P14, P15). By Theorem 6 the charm channel is the selected first-gap witness. Hence any admissible claim has the factorised form; a charm Yukawa introduced without it has no derived VERSF origin.

Gate-sufficiency. Conversely, if the factorised bridge is supplied with all premises satisfied, all normalisations declared, and $\delta\_\chi{,}_1$ fixed independently of charm, then within the FH-1 audit class the charm channel is admitted as the first-gap candidate: its value arises from the chain carrier → first gap → sector bridge → scheme-declared Yukawa rather than by insertion.

Limits of sufficiency. Gate-sufficiency does not entail full physical admissibility. Even with the bridge satisfied, the flavour sector can fail through anomaly inconsistency, incorrect chirality or representation content, unitarity violation, incompatible mixing, unstable running, nonlocal or nonfactorisable corrections, or failure against related observables. Those checks belong to other gates and to the programme audit (D11); FH-1 asserts admission to candidacy, not phenomenological completion. ■

17.3 Corollary — Calibration is not prediction

A paper that chooses $\delta\_\chi_{,1}$, $\mathscr{F}\_\chi$, or $\mathcal{N}_u$ to reproduce charm has calibrated charm, not predicted it; the formal blockade is Theorem 10.

---

18. Theorem 8 — Rare-Decay/Mass-Hierarchy Separation Theorem

18.1 Statement

Theorem 8 — Rare-Decay/Mass-Hierarchy Separation Theorem. Within FH-1, charm-hinge mass support (U_*, $\delta\_\chi_{,1}$) and charm-loop support in rare $b \to s\ \mu^+\mu^-$ amplitudes are distinct typed objects with no admissible identification. Consequently (i) no rare-decay observation — including the $B^0 \to K^{*0}\ \mu^+\mu^-$ tensions recorded in §10″ — constitutes evidence for or against FH-1, and (ii) no FH-1 result constitutes control over charm-loop amplitudes, $b \to s$ access, muon-current structure, or Wilson-coefficient shifts such as ΔC_9^μ, unless and until the D12 effective-operator bridge is derived.

Depends on: P4, P17, P19, I4, I10.

18.2 Proof

By P4, the hinge object (U_*, $\delta\_\chi_{,1}$) is typed as mass-eigenchannel support on $\mathcal{U}$_phys, related to mass only through the Theorem 3 bridge. Charm-loop support is not an element of $\mathcal{U}$_phys: it is a nonlocal amplitude contribution inside a loop-level flavour-changing neutral-current transition, typed by hadronic dynamics, effective-operator structure, and lepton-current access. By P17 mass support and transition support are already distinct classes; by P19 no identification between the two charm objects is admissible within FH-1. Inference in either direction would require a stated map between the classes; by I4 such a map is a bridge, and no such bridge is derived here — it is the open debt D12. Inference through an underived bridge is inadmissible, giving (i) and (ii). By I10, any future discharge of D12 must additionally treat CP-averaged and CP-asymmetry observables as separate diagnostics, since the observed pattern is asymmetric between them. ■

18.3 Corollary — Stress test, not verdict; vocabulary discipline

The $B^0 \to K^{*0}\ \mu^+\mu^-$ tensions are logged for the successor FCNC gate, whose target is the effective-operator deformation C_9^μ → C_9^μ + ΔC_9^μ firewalled against charm-loop uncertainty. Downstream papers must write "charm hinge"/"U_*" for the mass object and "charm-loop support" for the transition object, never unqualified "charm" where both readings are live.

19. Theorem 9 — Hinge Stability Theorem

19.1 Statement

Theorem 9 — Hinge Stability Theorem. Let $\hat{\chi}$_phys possess ordered nondegenerate eigenvalues $\lambda_0 < \lambda_1 < \lambda_2$. Under any admissible perturbation $\delta\hat{\chi}$ satisfying

$$\|\delta\hat{\chi}\| < \tfrac{1}{2} \cdot \min\{ \lambda_1 - \lambda_0 , \lambda_2 - \lambda_1 \} ,$$

the eigenvalue ordering remains unchanged and the first lifted eigenchannel remains continuously connected to U_*. Under exact χ/mass alignment ($[\hat{\chi}, Y_u{}^\dagger Y_u] = 0$), the charm-hinge identification is therefore not destroyed by sub-gap radiative, threshold, or bridge corrections; the approximate-alignment transfer is deferred to D13.

Depends on: P6, P10, P17A.

19.2 Proof

Let $g = \min\{\lambda_1 - \lambda_0, \lambda_2 - \lambda_1\}$ and let $\delta\hat{\chi}$ be an admissible self-adjoint perturbation with $\|\delta\hat{\chi}\| < g/2$. By Weyl's inequality [1, 2] each perturbed eigenvalue λ_i' of $\hat{\chi}$_phys + $\delta\hat{\chi}$ satisfies $|\lambda_i' - \lambda_i| \leq \|\delta\hat{\chi}\| < g/2$. The perturbed eigenvalues therefore lie in disjoint intervals of radius g/2 centred on λ_0, λ_1, λ_2, so the ordering $\lambda_0' < \lambda_1' < \lambda_2'$ is preserved and no eigenvalue crossing occurs. Along the continuous family $\hat{\chi}$_phys + $t\delta\hat{\chi}$, $t \in [0, 1]$, the spectral projections vary continuously and, absent crossings, the first lifted eigenspace at $t = 1$ is continuously connected to the eigenspace of λ_1 at $t = 0$ — that is, to U_*. For exact alignment ($[\hat{\chi}, Y_u{}^\dagger Y_u] = 0$) the identification of that eigenspace with the charm mass eigenchannel is likewise stable, since the perturbation cannot move the first gap to a different linear combination while remaining below the bound. For approximate alignment (ε_align > 0) the transfer requires a combined bound — perturbation norm and alignment slack jointly below the half-gap — whose derivation belongs to the alignment debt D13; the mass-eigenchannel conclusion is asserted here for the exact-alignment case only. ■

19.3 Corollary — Robust hinge, quantified

The charm-hinge identification tolerates all admissible corrections of operator norm below half the smaller adjacent spectral gap (exact alignment; the approximate-alignment transfer is deferred to D13); any claimed correction above that bound must be audited as a potential hinge reassignment (F37). The $\Gamma \geq \Gamma$_min condition of P10 guarantees strictly positive adjacent gaps, so the bound is nonvacuous.

20. Theorem 10 — No-Parameter-Hiding Theorem

20.1 Statement

Theorem 10 — No-Parameter-Hiding Theorem. Let the predicted charm quantity depend on parameters θ, partitioned as

$$\theta = \theta\_prior \cup \theta\_FH1 \cup \theta\_empirical \,,$$

where θ_prior are inherited from prior gates, θ_FH1 are fixed within FH-1's premised structure, and θ_empirical are external calibration inputs. Define the charm-proxy class

$\mathcal{O}$_charm = { observables whose values are dominantly controlled by charm-sector masses } ⊇ { y_C^obs , m_c^obs , the J/ψ and charmonium spectrum , D-meson masses , open-charm cross-sections } ,

and require the widened disjointness

$$\theta\_empirical \cap \mathcal{O}\_charm = \emptyset$$

during derivation, together with a declared charm-insensitivity bound on every component of the fixing set,

$$|\, \partial \ln \mathcal{O} \,/\, \partial \ln m\_c \,| \le \kappa\_c \text{ for every } \mathcal{O} \in \mathcal{O}\_prior \,,$$

with κ_c pre-declared, and the local identifiability condition

$$\mathrm{rank}(\, \partial\mathcal{O}\_prior \,/\, \partial\theta\_FH1 \,) = \dim \theta\_FH1 \,,$$

where $\mathcal{O}$_prior is the vector of charm-insensitive observables and prior-gate constraints. Then every FH-1 parameter is fixed by charm-insensitive inputs before charm is evaluated, and no factor of the bridge can function as a hidden charm fit — direct or by proxy.

Depends on: P9, P16, P18.

20.2 Proof

Suppose some factor of the bridge — a normalisation constant, a parameter of $\mathscr{F}$_χ, or a threshold — functioned as a hidden charm fit. Then one of three cases holds. (a) An observed charm datum entered its determination directly, contradicting the widened disjointness θ_empirical ∩ $\mathcal{O}$_charm = ∅ and P16. (b) A charm datum entered by proxy — the parameter was fixed against an observable dominated by m_c, such as the J/ψ mass — in which case either that observable belongs to $\mathcal{O}$_charm, again contradicting the disjointness, or it was admitted into $\mathcal{O}$_prior with $|\partial \ln \mathcal{O}/\partial \ln m_c| > \kappa_c$, contradicting the declared insensitivity bound. (c) The parameter was not determined by any input at all, i.e. it remained a free direction. In case (c), a free direction in θ_FH1 is a nonzero vector in the kernel of $\partial\mathcal{O}$_prior/∂θ_FH1, so

$$\mathrm{rank}(\partial\mathcal{O}_{prior} / \partial\theta_{FH1}) < \dim \theta_{FH1} ,$$

contradicting the identifiability condition. Hence under the three conditions every θ_FH1 component is locally pinned by charm-insensitive inputs, and by P18 none may be re-tuned after charm inspection. The bridge therefore contains no degree of freedom through which the observed charm value can be smuggled in, whether directly or through a charm-dominated proxy. ■

20.3 Corollary — Non-insertion is checkable

Non-insertion is no longer only a discipline: it is a verifiable rank condition on the parameter Jacobian plus a declared charm-insensitivity bound on the fixing set — auditable before any charm comparison, and closed against proxy smuggling through charmonium, D-meson, or open-charm data (F39).

The rank condition establishes **local** identifiability only: it pins θ_FH1 in a neighbourhood and does not exclude multiple disconnected solutions consistent with $\mathcal{O}$_prior. Global uniqueness is a separate, stronger claim not made here.

20.4 Frozen blind-prediction protocol

Theorem 10 becomes an experimentally auditable blind-prediction protocol under the following discipline:

1. Publish the parameter set θ and its provenance table (which observable fixed which parameter).
2. Freeze all FH-1 parameters before accessing any charm datum or $\mathcal{O}$_charm proxy.
3. Hash or timestamp the frozen parameter file.
4. Perform the charm comparison only afterwards.
5. Report every prior observable used, with its declared charm-sensitivity bound $|\partial \ln \mathcal{O}/\partial \ln m_c| \le \kappa_c$.

Any deviation from this order is an audit violation under P18 and F39.

---

21. Falsification Conditions

#	Failure	Consequence
F1	No Higgs-radial or equivalent mass carrier is supplied	No mass-support bridge
F2	No χ readout or operator is defined	No χ-gap theorem
F3	χ has no invariant ordering	No ordered spectrum
F4	χ is defined as observed mass	Circular derivation
F5	Up-type channels are χ-degenerate	No χ hierarchy
F6	Positive spectrum accumulates at the floor; no least separation exists	No first gap
F7	No χ metric, spectral distance, or access functional is supplied	No magnitude theorem
F8	$\delta_{\chi,1} = 0$ while hinge separation is claimed	Degenerate hinge failure
F9	χ magnitude depends on labels, coordinates, or conventions	Frame/label failure
F10	Bridge is not oriented-monotone and no suppression factor is named	Bridge inconsistency
F11	Carrier, χ magnitude, sector normalisation, and running are not separated	Hidden-fit failure
F12	Discriminators Φ, Γ, Σ or the selection rule are not defined	No hinge-selection structure
F13	Floor channel is used as clean first-gap witness while light-floor debts remain open	Floor failure
F14	Endpoint channel is used as first-gap witness while saturation is claimed	Endpoint failure
F15	U_* is named charm without the two-stage bridge — sector by representation content, channel by alignment and order transfer — or mass ordering is used in selection rather than at the definitional naming step, or a charm-specific representation datum is claimed	Naming failure
F16	Charm mass is claimed without scale and scheme	Scheme failure
F17	QCD running is absorbed into χ	Dressing/χ confusion
F18	Confinement effects are ignored in numerical comparison	QCD bridge failure
F19	CKM mixing is used without a mixing bridge	Mixing failure
F20	CKM closure is claimed from FH-1 alone	Scope violation
F21	$\delta_{\chi,1}$ is fitted from charm	Retrodiction failure
F22	$\mathcal{F}_\chi$ is fitted from charm	Hidden Yukawa failure
F23	$\mathcal{N}_u$ is chosen separately for charm	Sector-normalisation failure
F24	$\mathcal{D}_C$ becomes an arbitrary fit factor	Running-factor failure
F25	The up mass is claimed from FH-1 alone	Scope violation
F26	The top mass is claimed from FH-1 alone	Scope violation

#	Failure	Consequence
F27	Full up/charm/top hierarchy is claimed from FH-1 alone	Scope violation
F28	Down-type or lepton analogues are claimed without sector bridges	Cross-sector overreach
F29	Numerical charm prediction is claimed without D1–D7 discharged	Numerical overclaim
F30	Passing FH-1 is treated as empirical confirmation	Truth/admissibility confusion
F31	Channel roles are assigned after seeing observed masses	Retrodictive labelling failure
F32	χ gaps are counted twice as both hierarchy and running effects	Double-counting failure
F33	Charm hinge support is confused with charm-loop hadronic contamination	Category error
F34	$B^0 \to K^{*0} \mu^+\mu^-$ tensions are claimed as evidence for FH-1 without the FCNC bridge	Empirical overreach
F35	A Wilson-coefficient shift such as ΔC_9^μ is inferred from χ structure without the operator bridge	Operator-bridge failure
F36	CP-averaged and CP-asymmetry observables are treated as one diagnostic	Observable-class failure
F37	An admissible basis rotation or super-bound perturbation moves the first gap to another linear combination	Alignment/stability failure
F38	Discriminator thresholds Φ_max, Σ_max, Γ_min are tuned after inspecting the mass spectrum	Retrodictive-threshold failure
F39	The identifiability rank condition fails, a free FH-1 parameter direction remains, or an FH-1 parameter is fixed against a charm-proxy observable (charmonium spectrum, D-meson masses, open-charm cross-sections) or any observable exceeding the declared charm-insensitivity bound	Parameter-hiding failure
F40	Access magnitude and obstruction cost are conflated in the bridge orientation	Orientation failure
F41	An admissible substrate stratum exists unoccupied between the floor and the hinge, or occupancy is assumed without P6A	Occupancy failure
F42	One channel satisfies both $\Phi \geq \Phi$_max and $\Sigma \geq \Sigma$_max, or the interior channel fails $\Gamma \geq \Gamma$_min, emptying the feasible set	Discriminator-exclusivity failure
F43	δ_χ,₁ is treated as convention-independent while the monotone spectral-recalibration freedom of $\hat{\chi}$ remains unfixed	Recalibration failure
F44	A nonfactorisable operator-valued correction is absorbed into a scalar dressing factor despite mixing χ channels or threatening the mass ordering, or $\lVert \Delta Y^{nf} \rVert$ meets or exceeds ε_nf	Nonfactorisable-remainder failure
F45	An adjacent non-χ log-ratio exceeds the corresponding log-f_χ gap, reversing the mass ordering, so the first-lifted χ channel is not the intermediate-mass channel	Order-preservation failure

22. FH-1 Closure Certificate

FH-1 condition	Result
Higgs-radial carrier inherited	Inherited
Need for flavour structure beyond the carrier stated	Closed
χ realised as positive self-adjoint spectral operator	Premised (P6)
Well-founded first gap $\delta\_\chi,_1$ exists	Conditionally proved
Level II displacement architecture required and typed	Closed; physical calibration and numerical value open (D2, D3)
First occupied gap (P6 alone)	Conditionally closed
First admissible gap (P6 + P6A)	Conditional on P6A; SC-2 discharge unverified (D14)
Three closure levels distinguished (ordinal/metric/dynamical)	Closed
Oriented monotonicity of access bridge declared	Closed
Factorised bridge grounded at operator level	Conditionally proved (P9)
Dimensional consistency of bridge declared	Closed
Floor exclusion by discriminator Φ	Conditionally proved
Endpoint exclusion by discriminator Σ	Conditionally proved
Unique hinge selected without mass data	Conditionally proved
Two-stage charm identification (P13)	Conditionally proved; Stage 2 naming conditional on P8A (Theorem 6A) and exact alignment (P17A)
χ/mass alignment at gate order	Premised (P17A); derivation Open (D13)
Lowest-strata occupancy of $\mathcal{X}$_phys	Premised (P6A, SC-2 inheritance); derivation Open (D14)
Occupancy identity $\chi(U_*)$ = min admissible upper set	Conditionally proved (Lemma 1)
Hinge stability under sub-gap perturbations	Conditionally proved (Theorem 9)
Non-insertion enforced by parameter partition and rank condition	Closed (Theorem 10)
Scheme-declared, dimensionless Yukawa target	Closed
Running/dressing separated	Closed
Rare-decay/mass-hierarchy type separation	Closed (Theorem 8)

FH-1 condition	Result
Bridge-debt ledger D1–D15 recorded	Closed
Numerical $\delta_{\chi,1}$; access-to-mass law; charm Yukawa; charm mass	Open (D1–D7)
Up mass; top mass; full hierarchy; CKM; lepton sectors	Outside scope
$b \to s\, \mu^+\mu^-$ operators, charm loops, Wilson shifts	Outside scope (D12)
Empirical confirmation	Outside scope

FH-1 closure grade: Level I conditionally closed under P2–P6, with P6A additionally where "admissible" is claimed; Level II displacement architecture closed, physical calibration and numerical value open (D2, D3); Level III open (D1–D7); charm identification conditional on net mass-order preservation (P8A, Theorem 6A) and exact alignment (P17A); rare-decay typing closed.

---

23. Conclusion

Under the explicitly stated χ-spectrum, channel-count, floor, endpoint, alignment, and identification premises, FH-1 closes the first flavour-magnitude gate at the level the stated premises support — and states precisely which levels remain open.

The Higgs-radial theorem supplies a mass carrier; a carrier is not a hierarchy. FH-1 sharpens χ into a positive spectral operator whose well-founded spectrum yields a first physical gap

$$\delta_{\chi,1} = \lambda_1 - \lambda_0 > 0 \,,$$

an invariant spectral object rather than a verbal ladder step. Three closure levels are kept separate: the ordering (conditionally closed), the invariant magnitude (structure closed, value open), and the numerical Yukawa (architecture built, value open pending D1–D7).

The hinge is not labelled; it is selected. Substrate discriminators Φ, Γ, Σ with pre-declared thresholds fix a unique interior first-lifted channel U_* from spectral structure alone, and only then does the two-stage bridge name it: representation content fixes the up-type sector, and the P17A pairing with the Theorem 6A (P8A) order transfer attaches U_* to the intermediate-mass eigenchannel — the Standard Model's definition of charm — with mass ordering entering solely at that naming step, the identification made basis-invariant by the alignment premise and proved stable under all admissible perturbations below half the smaller adjacent spectral gap. The Yukawa may enter only through the dimensionless factorised bridge

$$y_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathscr{R}_H(\mu) \cdot \mathscr{F}_\chi(\delta_{\chi,1}) \cdot \mathcal{D}_C(\mu) \,,$$

grounded at operator level in a block-diagonal, multiplicatively renormalised effective action, with non-insertion enforced not merely as discipline but as a checkable parameter-partition and rank condition. Charm's separate appearance as a loop contribution in rare $b \to s\,\mu^+\mu^-$ decays is proved to be a different typed object, with all cross-inference deferred to the FCNC gate.

The programme sequence is:

HR-1 (Higgs-radial carrier) → FH-1 (first χ gap / hinge selection / rare-decay separation) → FCNC-1 ($b \to s\,\mu^+\mu^-$ closure-loop and muon-access shift) → CKM / mixing / lepton-access completion, seeded by the controlled failure of the P17A alignment (D13) .

Thus the charm channel is not an imported Standard Model Yukawa and not a generational label. It is the minimal admissible non-floor up-type channel within the three-state χ spectrum — the first metrically nonzero step of ordered flavour access.

References

[1] H. Weyl, "Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen," Math. Ann. 71 (1912) 441–479. Eigenvalue perturbation inequality used in Theorem 9.

[2] R. Bhatia, *Matrix Analysis*, Graduate Texts in Mathematics 169, Springer (1997). Weyl and eigenvector perturbation bounds.

[3] C. Davis and W. M. Kahan, "The rotation of eigenvectors by a perturbation. III," SIAM J. Numer. Anal. 7 (1970) 1–46. Spectral-projector ($\sin\theta$) bounds invoked in P17A.

[4] S. A. Gershgorin, "Über die Abgrenzung der Eigenwerte einer Matrix," Izv. Akad. Nauk SSSR, Otd. Mat. Est. Nauk, no. 6 (1931) 749–754. Disc separation used in Appendix B.

[5] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, revised edition, Academic Press (1980). Self-adjoint spectral theory and simultaneous diagonalisability.

[6] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*, Westview Press (1995). Standard Model representation content, Yukawa sector, renormalisation conventions.

[7] Particle Data Group, S. Navas et al., "Review of Particle Physics," Phys. Rev. D 110 (2024) 030001. Quark masses, scheme dependence, and naming conventions.

[8] LHCb Collaboration, R. Aaij et al., "A comprehensive analysis of the $B^0 \to K^{*0}\mu^+\mu^-$ decay," arXiv:2512.18053, submitted to Phys. Rev. Lett. (LHCb-PAPER-2025-041). The 8.4 fb^{-1} angular

record of §10″: full CP-averaged and CP-asymmetric observable sets, first S-wave observable set, muon-mass effects.

[9] CMS Collaboration, A. Hayrapetyan et al., "Angular analysis of the $B^0 \to K^*(892)^0\mu^+\mu^-$ decay in proton–proton collisions at $\sqrt{s} = 13$ TeV," Phys. Lett. B 864 (2025) 139406; arXiv:2411.11820. The 140 fb^{-1} record of §10″.

[10] LHCb Collaboration, R. Aaij et al., "Test of lepton universality in $b \to s\ell^+\ell^-$ decays," Phys. Rev. Lett. 131 (2023) 051803, arXiv:2212.09152; and "Measurement of lepton universality parameters in $B^+ \to K^+\ell^+\ell^-$ and $B^0 \to K^{*0}\ell^+\ell^-$ *decays," Phys. Rev. D 108 (2023) 032002, arXiv:2212.09153. R_K and R_K*.

[11] G. Buchalla, A. J. Buras and M. E. Lautenbacher, "Weak decays beyond leading logarithms," Rev. Mod. Phys. 68 (1996) 1125–1244. Effective weak Hamiltonian and Wilson-coefficient (C_9, C_{10}) conventions.

[12] LHCb Collaboration, R. Aaij et al., "Comprehensive analysis of local and nonlocal amplitudes in the $B^0 \to K^{*0}\mu^+\mu^-$ decay," JHEP 09 (2024) 026, Erratum ibid. 05 (2025) 208; arXiv:2405.17347. The unbinned local/nonlocal amplitude analysis over the full dimuon spectrum without charmonium vetoes — the experimental state of the art on the charm-loop and nonlocal-amplitude object that debt D12 targets.

Appendix A — Minimal Algebraic Skeleton

Channel space and spectral operator:

$\mathcal{U}$_phys three-dimensional (U_F = Φ-excluded , U_T = Σ-excluded , U_* = interior) , $\hat{\chi}$_phys = P_phys · $\hat{\chi}$ · P_phys , Spec_phys($\hat{\chi}$) ∋ $\lambda_0 < \lambda_1 < \lambda_2$.

Well-founded first gap:

$S_+(\lambda_0) = \{ \lambda \in \mathrm{Spec\_phys}(\hat{\chi}) : \lambda > \lambda_0 \} \neq \emptyset$, $\lambda_1 = \min S_+(\lambda_0)$, $\delta\_\chi{,}_1 = \lambda_1 - \lambda_0 > 0$.

Admissible support set and occupancy:

$\mathcal{X}$_phys ⊆ $\mathcal{X}$; Spec_phys($\hat{\chi}$) = χ($\mathcal{U}$_phys) ⊆ $\mathcal{X}$_phys ; P6A: Spec_phys($\hat{\chi}$) = the three ⪯_χ-least elements of $\mathcal{X}$_phys ; Lemma 1: $\chi(U_*) = \lambda_1 = \min \{ \chi \in \mathcal{X}\_phys : \chi_0 \prec\_\chi\, \chi \}$.

Displacement and magnitude:

$\Delta\_\chi(U\_F, U_*)$, $\delta\_\chi{,}_1 = \|\Delta\_\chi(U\_F, U_*)\|\_\chi = d\_\chi(\chi_0, \chi_1)$.

Discriminators and selection:

$\Phi(U), \Gamma(U), \Sigma(U)$; $U_* = \mathrm{argmin}\ \{ d\_\chi(\chi(U), \chi_0) > 0 : \Phi(U) < \Phi\_max , \Sigma(U) < \Sigma\_max , \Gamma(U) \geq \Gamma\_min \}$.

Two-stage identification bridge:

Stage 1 (sector): $B\_SM(U_i) = [\ SU(3)\_c , SU(2)\_L , Y , Q , \text{chirality}\ ]_i$; $(3, 2, +1/6)\_L \oplus (3, 1, +2/3)\_R$. Stage 2 (channel): P17A pairing + Theorem 6A (P8A) order transfer ; $U_* \mapsto$ intermediate-mass eigenchannel $\stackrel{\text{def}}{=}$ c .

Alignment:

$[\ \hat{\chi}, Y_u \dagger Y_u\ ] = 0$ or $\| [\ \hat{\chi}, Y_u \dagger Y_u\ ] \| \leq \varepsilon\_align$ (exact case consumed by FH-1 identification statements; approximate case admissible only under the Davis–Kahan projector bound, deferred to D13).

Operator bridge and projection:

$\mathscr{L}\_eff \supset - \bar{Q}\_L \cdot Y_u(\mu) \cdot \tilde{H} \cdot U\_R + h.c.$; $Y_u(\mu) = Z\_L\text{^}(-1/2)(\mu) \cdot B_u \cdot f\_\chi(\hat{\chi}) \cdot Z\_R\text{^}(-1/2)(\mu)$; $y_i(\mu) = Z\_L,_i\text{^}(-1/2)(\mu) \cdot B_{u,i} \cdot f\_\chi(\lambda_i) \cdot Z\_R,_i\text{^}(-1/2)(\mu)$.

Scalar bridge (all factors dimensionless):

$y\_C(\mu) = \mathcal{N}_u(\mu) \cdot \mathscr{R}\_H(\mu) \cdot \mathscr{F}\_\chi(\delta\_\chi,_1) \cdot \mathcal{D}\_C(\mu)$, $\mathscr{R}\_H(\mu) = g\_H,U(\mu)/g\_H,ref(\mu)$, $m\_C(\mu) = (v(\mu)/\sqrt{2}) \cdot y\_C(\mu)$.

Reduced support and closure levels:

$\hat{y}\_C(\mu) = \mathscr{F}\_\chi(\delta\_\chi,_1)$; Level III target: $y\_C(\mu)/y\_ref(\mu) = \boldsymbol{\mathcal{M}}\_\chi(\delta\_\chi,_1 ; \lambda\_substrate)$.

Orientation:

$d\mathscr{F}\_\chi/d\delta\_\chi > 0$, $d\mathscr{F}\_\chi/d\mathcal{A}\_\chi < 0$.

Stability bound:

$\|\delta\hat{\chi}\| < \frac{1}{2} \cdot \min\{ \lambda_1 - \lambda_0 , \lambda_2 - \lambda_1 \}$.

Parameter partition and identifiability:

$\theta = \theta\_prior \cup \theta\_FH1 \cup \theta\_empirical$; $\theta\_empirical \cap \mathcal{O}\_charm = \emptyset$; $|\partial \ln \mathcal{O}/\partial \ln m\_c| \leq \kappa\_c\ \forall\ \mathcal{O} \in \mathcal{O}\_prior$; $\mathrm{rank}(\ \partial\mathcal{O}\_prior/\partial\theta\_FH1\ ) = \dim \theta\_FH1$.

Rare-decay separation targets (D12 only):

$\mathcal{O}_9\text{^}\mu = (\bar{s}\ \gamma\_\alpha\ P\_L\ b)(\bar{\mu}\ \gamma\text{^}\alpha\ \mu)$, $\mathcal{O}_{10}\text{^}\mu = (\bar{s}\ \gamma\_\alpha\ P\_L\ b)(\bar{\mu}\ \gamma\text{^}\alpha\ \gamma_5\ \mu)$, $C_9\text{^}\mu \to C_9\text{^}\mu + \Delta C_9\text{^}\mu$.

Appendix B — Worked 3×3 Spectral Toy Model

This appendix illustrates the FH-1 audit with an explicit spectral toy. It is not a numerical charm derivation.

Let the χ operator on $\mathcal{U}$_phys be, in a weak-channel basis,

$$\hat{\chi} = \begin{pmatrix} c & \varepsilon_{12} & 0 \\ \varepsilon_{12}^* & c+a & \varepsilon_{23} \\ 0 & \varepsilon_{23}^* & c+b \end{pmatrix} , \quad 0 < a < b ,$$

with off-diagonal couplings ε_{12}, ε_{23} small compared with the diagonal separations. The offset c secures positivity of $\hat{\chi}$ — the unshifted diagonal would let the lowest eigenvalue dip below zero at second order — and, since spectral gaps are offset-invariant, changes nothing in the gap and stability analysis. Project into the physical sector and diagonalise.

Exact ordering (Gershgorin [4]). Sufficient conditions for three positive, pairwise-disjoint spectral intervals with exactly one eigenvalue in each are

$$c > |\varepsilon_{12}| , \quad a > 2|\varepsilon_{12}| + |\varepsilon_{23}| , \quad b - a > |\varepsilon_{12}| + 2|\varepsilon_{23}| :$$

under these, the three Gershgorin discs are disjoint and ordered, so the spectrum is exactly $\lambda_0 < \lambda_1 < \lambda_2$ with $\lambda_0 > 0$ — an exact statement, not a perturbative one. The second-order formulas below are approximations within this exactly established ordering; they are not the ordering's proof.

Eigenvalue structure. To second order in the couplings,

$$\lambda_0 \approx c - |\varepsilon_{12}|^2/a > 0 , \quad \lambda_1 \approx c + a + |\varepsilon_{12}|^2/a - |\varepsilon_{23}|^2/(b - a) , \quad \lambda_2 \approx c + b + |\varepsilon_{23}|^2/(b - a) .$$

Under the explicit inequalities

$$|\varepsilon_{12}|^2 < a^2/4 \text{ and } |\varepsilon_{23}|^2 < a(b - a)/4 ,$$

the ordering $\lambda_0 < \lambda_1 < \lambda_2$ — already exact by the Gershgorin conditions — persists in the second-order formulas, the positive upper set above λ_0 has least element λ_1, and the first gap $\delta\_\chi{,}_1 = \lambda_1 - \lambda_0 \approx a + 2|\varepsilon_{12}|^2/a - |\varepsilon_{23}|^2/(b - a)$ — offset-independent, as gaps must be — is positive and well defined. The lowest eigenstate is floor-dominated (large Φ: its stratum is compressed toward the boundary and its composition is dominated by the floor basis vector); the highest is endpoint-dominated (large Σ: dominated by the b stratum); the middle eigenstate satisfies neither condition and satisfies $\Gamma \geq \Gamma$_min, since its stratum is separated from both neighbours by finite gaps.

Selection. The rule of §3.3 selects the middle eigenstate as U_* with $\chi(U_*) = \lambda_1$ — using only the spectrum and the discriminators, with no mass input.

Occupancy. The toy takes $\mathcal{X}$_phys equal to the realised spectrum, so P6A holds trivially and Lemma 1 is automatic. A genuine occupancy test requires a substrate support lattice strictly larger than the spectrum — exactly what debt D14 sends to the SC-2 census machinery.

Stability. Any admissible perturbation $\delta\hat{\chi}$ with

$$\|\delta\hat{\chi}\| < \tfrac{1}{2} \cdot \min\{\, \lambda_1 - \lambda_0 \,,\, \lambda_2 - \lambda_1 \,\}$$

preserves the ordering and keeps the first lifted eigenchannel continuously connected to U_* (Theorem 9). The hinge identity in the toy is therefore robust: it cannot be changed by small radiative, threshold, or bridge corrections, and any correction that could change it is detectably super-bound.

Bridge. With an illustrative increasing access law $\mathscr{F}$_χ(δ) — for example $\mathscr{F}$_χ(δ) = e^(−A/δ) with constant A > 0, increasing in δ per P8 — the hinge's reduced support is ŷ_∗ = $\mathscr{F}$_χ(δ_χ,₁), while the floor and endpoint supports are governed by their own functionals with floor-suppression and saturation content respectively. The toy exhibits the logical roles — floor, selected hinge, endpoint — and the stability bound; it does not exhibit observed values.

Appendix C — χ-Gap Audit Table

Object	Meaning	Admissible role	Failure if confused
$\mathcal{U}$_phys	Physical up-type channel space	Domain of the audit	Undefined channel set
$\hat{\chi}$_phys	Positive self-adjoint χ operator	Carries ordered access spectrum	Mass renamed as χ
$\lambda_0, \lambda_1, \lambda_2$	Physical spectrum	Floor, first lifted, endpoint values	Retrodictive labelling
$S_+(\lambda_0)$	Positive upper set	Well-foundedness carrier	Accumulation at floor ignored
$\mathcal{X}$_phys	Substrate-admissible support set	Domain of the admissible minimum	Conflated with occupied spectrum (F41)
$\delta\_\chi,_1$	First physical gap	Hinge-magnitude source	Fitted from charm
Δ_χ(U, U′)	Typed displacement	Ordered stratum movement	Coordinate subtraction assumed
‖·‖_χ, d_χ	Invariant magnitude/distance	Metric closure carrier	Rank mistaken for magnitude
Φ(U)	Floor proximity	Excludes floor channel	Floor used as clean witness

Object	Meaning	Admissible role	Failure if confused
Γ(U)	Interior isolation	Selection-rule cleanliness clause; certifies nonvacuous Theorem 9 bound	Threshold tuned retrodictively (F38)
Σ(U)	Endpoint saturation	Excludes endpoint channel	Saturation counted as first gap
Φ_{max}, Σ_{max}	Pre-declared thresholds	Fix roles before mass data	Tuned retrodictively (F38)
U_*	Selected hinge	First-gap witness	Treated as fitted generation
$B_{SM}(U)$	Representation datum	Stage-1 sector fixing only	Claimed to select among generations (F15)
P17A pairing + Theorem 6A (P8A) transfer	Channel-naming stage	Stage-2 definitional naming	Mass ordering used in selection (F15)
$[\hat{\chi}, Y_u^\dagger Y_u]$	Alignment commutator	Basis-invariance carrier; CKM seed	Basis-dependent "charm χ value"
$\mathscr{F}_\chi$	Access-to-mass functional	Reduced Yukawa support	Hidden fitted Yukawa
$\mathcal{A}_\chi$	Obstruction cost	Oppositely oriented to δ_χ	Conflated with magnitude (F40)
$\mathscr{R}_H$	Normalised carrier ratio $g_{H,U}/g_{H,ref}$	Dimensionless carrier support	Treated as dimensionful VEV/stiffness
$\mathcal{N}_u$	Up-sector normalisation	Sector-level scaling	Chosen per channel
$\mathcal{D}_C$	Running/dressing factor	Scheme/QCD bridge	Absorbed into χ
$y_C(\mu)$, $m_C(\mu)$	Scheme Yukawa; mass via $v(\mu)/\sqrt{2}$	Numerical comparison targets	Bare mass claim
θ partition	Prior/FH-1/empirical parameters	Non-insertion enforcement	Free direction hides charm fit (F39)
Charm-loop support	Nonlocal $b \to s\, \mu^+\mu^-$ amplitude	FCNC-gate object (D12)	Confused with hinge support (F33)
$\mathcal{O}_9^\mu$, $\mathcal{O}_{10}^\mu$, ΔC_9^μ	Effective operators and shift	D12 bridge targets	Claimed as FH-1 content

Appendix D — Referee Objections and Replies

Objection 1 — You premise floor, hinge, and endpoint roles, eliminate two, and call the survivor charm. Where is the new physical derivation?

Reply. This is the correct objection to any architecture that merely premises three roles and eliminates two; the paper's structure answers it as follows. The three roles are defined by substrate discriminators Φ, Γ, Σ with pre-declared thresholds, making no reference to observed masses or generation names; the hinge is fixed by an explicit selection rule consuming only ($\hat{\chi}$, Φ, Γ, Σ); and the Standard Model name is attached afterwards by the two-stage bridge — representation content fixing the sector, alignment and order transfer fixing the channel, with mass ordering entering only as the definitional naming convention. The result is not "the middle channel is charm, so charm is the hinge" but "VERSF selects an interior, first-lifted channel without using mass data; the two-stage bridge (P13) identifies that selected state as charm." The residual conditionality — that the discriminators and thresholds are premised, not yet derived from substrate primitives — is stated openly (P10–P12) and is attackable there, which is where a referee should attack it.

Objection 2 — Does FH-1 derive the charm mass?

Reply. No, and the closure-level hierarchy of §3.8 makes the claim-type explicit: ordinal closure is conditionally proved, metric closure is specified, dynamical closure is architecture only. Numerical prediction requires D1–D7.

Objection 3 — Is the "first increment" just a verbal ladder step?

Reply. No. Under P6 it is a spectral gap $\delta_{\chi,1} = \lambda_1 - \lambda_0$ of a positive self-adjoint operator with well-founded spectrum. If the spectrum accumulated at the floor, the theorem would fail — falsifier F6 — and no regulator is smuggled in.

Objection 4 — What if an admissible stratum between λ_0 and λ_1 is unoccupied?

Reply. Then charm would witness a later increment, not the first — which is exactly why occupancy is a named premise rather than a silent assumption. P6A, the FH-1 restriction of the SC-2 world-to-framework occupancy map, states that the three physical channels occupy the three lowest admissible strata of $\mathcal{X}$_phys, and Lemma 1 converts that premise into the identity $\chi(U_*) = \min\{ \chi \in \mathcal{X}_{phys} : \chi_0 \prec_\chi \chi \}$. An admissible unoccupied stratum below the hinge is falsifier F41, and the substrate derivation of occupancy is debt D14. Absent P6A, the theorem target weakens to "charm witnesses the first occupied increment"; the paper prefers the stronger premised claim with its falsifier exposed.

Objection 5 — Doesn't $\Delta\chi_1 = \chi_1 - \chi_0$ assume a linear structure on χ?

Reply. The invariant distance d_χ is primary; $\Delta_\chi(U_F, U_*)$ is a typed displacement and $\delta_{\chi,1}$ its scalar magnitude. Subtraction appears only in the spectral representation, where it is legitimate, and $\Delta\chi_1$ is reserved for a coordinate representative where an affine chart exists.

Objection 6 — Isn't the four-factor bridge an ansatz?

Reply. At FH-1 order it is the χ-eigenchannel projection of a block-diagonal, multiplicatively renormalised effective Yukawa operator (P9, §3.6). Beyond that order, multiplicative diagonal corrections enter the dressing factor, while nonfactorisable channel-mixing corrections remain in the bounded operator remainder ΔY^{nf}; absorbed into a scalar factor, they trigger F44.

Objection 7 — Before diagonalisation, "charm" is basis-dependent. Which state is U_*?

Reply. U_* is declared as a χ eigenchannel, and P17A supplies alignment with the mass eigenchannel — exact commutation of $\hat{\chi}$ with $Y_u^\dagger Y_u$, or, in the approximate case, the declared Davis–Kahan projector bound of P17A, not a raw commutator norm alone. If a basis rotation could move the first gap to another combination, the identification would fail (F37). The controlled failure of alignment is the designed seed of the CKM gate (D13).

Objection 8 — Small corrections could reshuffle the spectrum and move the hinge.

Reply. Theorem 9 gives the quantitative bound: any admissible perturbation of norm below half the smaller adjacent gap preserves the ordering and keeps the first lifted eigenchannel continuously connected to U_*. Hinge reassignment requires a detectably super-bound correction.

Objection 9 — A hidden normalisation constant could still be tuned to charm.

Reply. Theorem 10 blocks this formally, including the proxy route: the parameter set is partitioned; the exclusion class $\mathcal{O}_{charm}$ covers not only the charm mass and Yukawa but observables dominated by them — the charmonium spectrum, D-meson masses, open-charm cross-sections; every fixing observable must satisfy a declared charm-insensitivity bound $|\partial \ln \mathcal{O}/\partial \ln m_c| \leq \kappa_c$; and the identifiability rank condition guarantees no free FH-1 parameter direction survives. A hidden fit now shows up as a checkable rank deficiency or a bound violation (F39).

Objection 10 — What if there are more than three up-type channels, or the ordering is inferred from masses?

Reply. The three-channel premise is explicit (P10); extended sectors require a separate audit. Retrodictive ordering or threshold-tuning is excluded by P18 and falsifiers F31, F38.

Objection 11 — You put charm at the centre of a mass ladder while charm loops are under scrutiny in the b → s $\mu^+\mu^-$ tensions. Are you ignoring that?

Reply. The two roles are typed and separated by theorem. The hinge is $(U_*, \delta_{\chi,1})$, a mass-eigenchannel object; charm-loop contributions are transition-class amplitude content and a major theoretical uncertainty in interpreting selected rare-decay observables. Theorem 8 proves no inference runs in either direction absent the D12 bridge, and I10 requires CP-averaged and CP-asymmetry observables — whose recorded behaviour differs (§10″) — to be treated as separate diagnostics. Claiming the tensions as evidence for FH-1 is falsifier F34.

Objection 12 — What is the strongest single result?

Reply. That the first flavour separation is a well-founded spectral gap; that its witness is selected by substrate discriminators before any mass datum or generation name enters; that the selection is basis-invariant and stable under all sub-gap perturbations; and that any Yukawa content must pass through a factorised, operator-grounded, rank-audited bridge that cannot conceal a charm fit.

Appendix E — Final Theorem

Conditional First-Gap Charm-Hinge Theorem. Let the physical up-type VERSF channel space be three-dimensional and carry a positive, invariant, self-adjoint χ-access operator with a well-founded nondegenerate spectrum $\lambda_0 < \lambda_1 < \lambda_2$. Suppose the lowest eigenspace satisfies an independently defined floor criterion, the highest satisfies an independently defined endpoint-saturation criterion, the intermediate eigenspace satisfies neither, the physical channels occupy the lowest admissible strata of the substrate support set $\mathcal{X}_{phys}$, and the χ and up-type mass-support operators are aligned at FH-1 order. Then the intermediate eigenspace is the unique first lifted physical χ channel, stable under all admissible perturbations of norm below half the minimum adjacent spectral gap. If the two-stage identification bridge — representation content fixing the up-type sector, alignment and order transfer pairing the intermediate eigenspace with the intermediate-mass eigenchannel, which the Standard Model names charm — maps that eigenspace accordingly, charm is the unique first-gap up-type hinge, with mass ordering entering only at that definitional naming step. Its reduced Yukawa may depend on $\lambda_1 - \lambda_0$ only through an independently derived access-to-mass map; no observed charm datum may enter the construction of that gap, map, or channel identification; and this hinge role admits no identification with charm-loop support in rare $b \to s\, \mu^{+}\mu^{-}$ transitions absent a derived effective-operator bridge. Within the restricted FH-1 audit class, satisfaction of these conditions is necessary and gate-sufficient for admitting charm as the first-gap candidate; it is not sufficient for full phenomenological validity of the VERSF flavour sector.