The Minimal Hermitian Lift and Frame-Rotation Firewall in VERSF

▲ Programme Milestone — Standard Model Spectral-Frame Series Gate SF-1 / Hermitian-Lift and Frame-Invariant Prediction Closure

Preventing basis choices, matrix entries, texture conventions, and arbitrary frame rotations from being mistaken for physical Standard Model predictions

Keith Taylor VERSF Theoretical Physics Programme — Standard Model Spectral-Frame Series Tier A frame-discipline gate paper, SF-1

General Reader Summary

The occupancy paper SC-2 fixed the seating chart for Standard Model matter inside VERSF. It said where each quark, lepton, chirality, weak component, colour branch, and generation layer sits in the framework.

Once every particle has an address, a new danger appears.

A later paper might write down a matrix and say:

"This matrix predicts the electron mass," "This entry predicts the strange quark," "This rotation predicts CKM," "This texture explains the hierarchy."

None of that is yet legitimate. A matrix is a coordinate table. Its entries depend on the frame chosen to write it. A different weak basis, flavour basis, or generation frame changes the entries without changing the physics. Before VERSF is allowed to predict masses, Yukawas, CKM, PMNS, or hierarchy exponents, it needs a firewall between coordinate conventions and physical claims.

This paper supplies that firewall.

The core rule is:

A raw closure matrix is not itself a physical prediction. The first admissible physical object is its minimal Hermitian lift.

For a closure map Γ, the minimal Hermitian lifts are:

$K_L = \Gamma \Gamma^\dagger$, $K_R = \Gamma^\dagger \Gamma$.

These objects are Hermitian, positive, basis-covariant, and have frame-invariant eigenvalues. Their eigenvalues are the squared singular values of Γ. After electroweak symmetry breaking, these are the squared Yukawa or mass scales, up to the common Higgs factor.

In plain terms:

You cannot predict physics from arbitrary matrix entries. You may predict only frame-invariant spectral data, or physically meaningful relative rotations between Hermitian lifts.

For example, the up-quark and down-quark closure maps each have a left Hermitian lift on the shared quark weak-generation frame. The CKM matrix is not the absolute rotation of either one. It is the **relative rotation** between the up-sector and down-sector left eigenframes — and even that relative rotation is physical only after its residual rephasing and ordering freedom has been quotiented away. The comparison is meaningful at all only because the up and down sectors hang from the same weak doublet, which locks their left frames together. This paper states that lock as a named rule rather than leaving it implicit.

Likewise, PMNS is not available in the minimal massless-neutrino ledger. It becomes meaningful only after a neutrino-mass extension supplies a neutrino Hermitian lift on the shared lepton weak-generation frame.

The paper's key claim is:

Given the SC-2 occupancy map, any admissible VERSF mass, hierarchy, Yukawa, CKM, PMNS, or CP prediction must pass through the minimal Hermitian lift and must target frame-invariant spectral or relative-frame data. Raw matrix entries, arbitrary texture zeros, unconstrained basis choices, and single-sector diagonalisation conventions are not physical predictions.

In one sentence:

SF-1 makes VERSF matrix predictions audit-safe: masses come from Hermitian-lift spectra, mixings come from relative Hermitian-lift frames, and arbitrary basis choices are firewalled away.

Scope Box — What SF-1 Closes and What It Does Not

Claim	Status
Need for post-occupancy operator discipline	Closed
Closure map Γ defined between occupied Ω-address bundles	Closed
Frame rotations defined	Closed
Doublet-locked left-frame group (HLF-4′)	Closed

Claim	Status
Shared generation factor $\mathcal{G}$_Q made explicit	Closed
Minimal Hermitian lift defined	Closed
Left lift K_L = Γ Γ†	Closed
Right lift K_R = Γ† Γ	Closed
Metric-adjoint version for non-trivial positive kinetic metrics	Closed
Polar decomposition and uniqueness of the positive part	Closed
Basis covariance of Hermitian lifts	Closed
Invariance of singular spectrum	Closed
Minimality and uniqueness of the quadratic lift, relative to declared structure	Closed
Block-Hermitian auxiliary firewall	Closed
Canonicalisation equivalence for non-trivial metrics	Closed
Raw matrix-entry firewall	Closed
Texture-zero firewall	Closed unless symmetry-fixed
Invariant-expressibility rule for structural constraints	Closed
Single-sector diagonal-basis firewall	Closed
Relative-frame origin of CKM	Closed
Relative-frame parameter count (angles and phases)	Closed
Commutator–quartet CP identity, derived in-corpus	Closed
PMNS extension-dependence	Closed
PMNS invariant parameter counts per extension type	Closed
Takagi structure of Majorana mass operators	Flagged for extension gates
Rectangular closure maps	Scoped to extension gates
Degeneracy firewall	Closed
Rephasing/permutation firewall	Closed
CP-invariant target discipline	Closed
Requirement that later predictions target Ω-addresses and Hermitian-lift invariants	Closed
Numerical fermion masses	Not closed
Yukawa eigenvalues	Not closed
CKM numerical entries	Not closed
PMNS numerical entries	Not closed
CP phase values	Not closed
Neutrino-mass mechanism	Not closed
Higgs vev or electroweak scale	Not closed
Gauge-coupling values	Not closed

SF-1 is a **matrix-admissibility and frame-invariance theorem**. It does not predict the values of masses or mixings. It proves what kind of object a legitimate prediction must target.

Abstract

SC-2 fixed the world-to-framework occupancy map Ω for minimal Standard Model matter in VERSF. Once matter fields have fixed framework addresses, later quantitative gates require operators acting on those occupied slots. However, closure matrices, Yukawa matrices, texture patterns, and diagonalising rotations are frame-dependent objects. Their raw entries change under arbitrary weak-basis, flavour-basis, and generation-frame rotations. Without a firewall, a later derivation could mistake a coordinate convention for a physical prediction.

This paper defines the **minimal Hermitian lift** of an admissible closure map

$$\Gamma_f : \mathscr{R}_f \to \mathscr{L}_f ,$$

where $\mathscr{R}_f$ and $\mathscr{L}_f$ are right-terminal and left-attached Ω-address bundles for a given closure channel. In canonical kinetic coordinates the two Hermitian lifts are

$$K_f^L = \Gamma_f \Gamma_f^\dagger , \quad K_f^R = \Gamma_f^\dagger \Gamma_f .$$

For non-trivial positive kinetic metrics, the dagger is replaced by the metric adjoint ‡. The lifts are self-adjoint, positive semidefinite, basis-covariant, and have the same non-zero spectrum. Their eigenvalues are the squared singular values of Γ_f, and the left lift is exactly the square of the unique positive factor in the polar decomposition of Γ_f.

The paper proves four theorems.

1. **Minimal Hermitian Lift Theorem:** the first admissible positive spectral object associated with a closure map is the quadratic Hermitian lift, unique up to the left/right choice and positive normalisation, and equal to the squared polar factor of the closure map.
2. **Frame-Rotation Firewall Theorem:** raw matrix entries, arbitrary texture zeros, diagonal-basis choices, and single-sector eigenvectors are not physical predictions unless fixed by additional symmetry or re-expressed as frame-invariant conditions.
3. **Relative-Frame Mixing Theorem:** CKM and PMNS-type data arise only from relative orientations of Hermitian lifts acting on a shared left weak-generation frame; a single closure sector has no absolute mixing prediction; and the physical relative-frame content, after quotienting residual rephasing and ordering freedom, carries exactly $n(n-1)/2$ angles and $(n-1)(n-2)/2$ phases for n non-degenerate generations.
4. **Prediction-Admissibility Theorem:** any later VERSF mass, hierarchy, Yukawa, CKM, PMNS, or CP claim must specify an Ω-address target and a frame-invariant Hermitian-lift observable, and must pass five explicit admissibility tests.

The admissible frame group is itself disciplined by a doublet-locking premise, HLF-4′: left bundles that are weak components of a single occupied doublet address carry one shared generation frame and admit only a common left rotation. This premise is load-bearing — were independent left rotations of weak partners admissible, the up-sector and down-sector lifts could be diagonalised simultaneously and CKM would dissolve into convention. The paper also proves, as a structural identity derived in-corpus, the commutator–quartet relation connecting det [K_u^Q, K_d^Q] to the rephasing-invariant quartet of the relative-frame matrix and the eigenvalue splittings of the two lifts. This fixes the invariant class that any later VERSF CP theorem must target, without deriving its numerical value.

SF-1 therefore closes the frame-rotation gate. It permits VERSF to predict spectra, eigenvalue ratios, singular values, trace/determinant combinations, commutator invariants, projective relative-frame data, and CP invariants. It forbids treating raw matrices, movable labels, texture choices, basis conventions, or unconstrained rotations as physical results.

Contents

---

0. Inheritance, Firewalls, and Aim

0.1 Inherited results

SF-1 inherits the following.

I1 — QN-1 κ-ledger discipline. Localization stiffness, probability exponents, reciprocal compliance, channel-log weights, gap scales, and response couplings are typed separately. No bare κ may be used as a generic fitting parameter.

I2 — QN-2 Standard Model convention ledger. The active Standard Model convention is

$$Q = T_3 + Y/2 .$$

The Higgs doublet is

$$H = (H^+, H^0)^T , Y_H = +1 ,$$

and

$$\tilde{H} = i\sigma_2 H^* , Y_{\tilde{H}} = -1 .$$

The charged-fermion closure routes are:

$$\bar{Q}_L \tilde{H} u_R , \bar{Q}_L H d_R , \bar{L}_L H e_R .$$

I3 — Physical-sector positivity. Physical matter slots carry positive-norm terminal support. Negative-norm leakage, ghost residue, gauge artifacts, and unphysical adjoint remnants do not occupy matter slots.

I4 — SC-1 species and generation census. The minimal chiral matter census contains three generation layers, each with

$$Q_L , u_R , d_R , L_L , e_R ,$$

with no sterile right-handed neutrino in the minimal census.

I5 — SC-2 world-to-framework occupancy map. Every minimal chiral Standard Model matter field occupies a fixed Ω-address. The address includes generation layer, species role, colour, weak representation, chirality, hypercharge, electric charge, Higgs closure route, anomaly role, and extension status.

I6 — Generation-layer anchoring. The layer coordinate $i = 1, 2, 3$ is inherited from SC-1. Representation data alone separate species roles within a generation, but not the three representation-identical generation copies.

I7 — Bridge discipline. A later VERSF prediction must specify its Ω-address target. A particle name alone is not an admissible bridge.

I8 — Doublet address factorisation (corollary of I5 and I6, recorded here for citation). A doublet Ω-address factorises as generation layer × weak-component label. This follows from I5 and I6: by I6 the three generation copies are representation-identical, so the layer coordinate varies over identical representation data, and by I5 the weak-component label is fixed occupancy data within each address. The occupied doublet bundle is therefore spanned by the independent (layer, component) pairs. This inheritance line is load-bearing for §3.5 and hence for §§8 and 11.

0.2 New problem inherited from SC-2

SC-2 fixed addresses.

But a fixed address does not yet define an admissible operator.

Once the ledger records

$u_R \mapsto \mathfrak{U}_i$, $d_R \mapsto \mathfrak{D}_i$, $e_R \mapsto \mathfrak{E}_i$,

a later paper must still answer:

- what kind of object acts on these slots;
- which part of that object is physical;
- which part is a coordinate choice;
- what survives weak-basis rotation;
- what counts as a mass prediction;
- what counts as a mixing prediction;
- when a matrix entry is meaningful;
- when a texture zero is physical;
- when a diagonal form is merely a convention.

SF-1 answers these questions.

0.3 Firewalls

This paper does **not** claim:

- that any fermion mass is numerically derived;
- that any Yukawa eigenvalue is numerically derived;
- that CKM entries are numerically derived;
- that PMNS entries are numerically derived;
- that CP phases are numerically derived;
- that neutrino masses are generated in the minimal ledger;
- that the Higgs vev is derived;
- that a raw matrix entry is physical;
- that a texture zero is physical without a symmetry or invariant reason;
- that a single-sector diagonal basis is predictive;
- that an arbitrary unitary rotation is observable;
- that eigenvectors inside a degenerate subspace are physical;
- that rephasing conventions are physical;
- that generation ordering is derived from the Hermitian lift.

It proves a narrower but essential result:

Given fixed Ω-addresses, the first admissible spectral object associated with a closure map is its minimal Hermitian lift. Physical predictions may target only Hermitian-lift spectra, invariant functions of those spectra, or relative-frame invariants between Hermitian lifts sharing a common weak-generation frame.

0.4 Aim

The aim is to close the frame loophole before any numerical mass or mixing programme begins.

SC-2 says:

"Here are the occupied matter slots."

SF-1 says:

"Here is the only kind of matrix object that may be treated as physically predictive on those slots."

0′. Predictive-Content Ledger

Object	Prior status	Status here
Ω-addresses	Closed in SC-2	Used as operator domains/codomains
Closure map Γ	Informal	Defined

Object	Prior status	Status here
Left/right occupied bundles	Implicit	Defined
Frame rotation	Potential ambiguity	Defined
Doublet-locked left frames	Silent assumption	Closed as HLF-4′
Shared generation factor 𝒢_Q	Implicit	Closed as lemma, §3.5
Block-Hermitian auxiliary 𝒟_Γ	Potential objection	Firewalled as auxiliary representation
Canonicalised map Γ̂	Absent	Defined, spectrum-equivalent to ‡-lifts
Raw closure matrix	Tempting prediction target	Firewalled
Hermitian lift	Absent	Defined
K_L = Γ Γ†	Absent	Closed
K_R = Γ† Γ	Absent	Closed
Polar decomposition of Γ	Absent	Closed
Positive kinetic metric	Implicit	Included
Metric adjoint ‡	Absent	Defined
Singular spectrum	Informal	Prediction-eligible
Matrix entries	Often misused	Not prediction-eligible unless symmetry-fixed
Texture zeros	Often misused	Not prediction-eligible unless invariantly expressible
Invariant expressibility rule	Absent	Closed
Diagonal basis	Often overclaimed	Not prediction-eligible alone
CKM	Later target	Typed as relative left-frame invariant
CKM parameter count	Informal	Closed as relative-frame count
PMNS	Later extension target	Requires neutrino Hermitian lift
PMNS parameter counts	Absent	Closed per extension type
CP phase	Later target	Typed as commutator/quartet invariant
Commutator–quartet identity	Absent	Closed, derived in Appendix A
Takagi/Majorana operator structure	Absent	Flagged for extension gates
Degenerate eigenframes	Potential ambiguity	Firewalled
Rephasing/permutation	Potential ambiguity	Firewalled
Later Yukawa gates	Vulnerable to basis drift	Must target Hermitian-lift observables
Later hierarchy gates	Vulnerable to label drift	Must target Ω-address spectra

1. The Post-Occupancy Operator Problem

1.1 Why SC-2 is not enough

SC-2 fixed the addresses of e_R, u_R, d_R, Q_L, and L_L.

But a mass theorem does not act on a name. It acts on a space.

The charged-lepton closure sector has

$\mathscr{L}_e$ for the left charged-lepton components,

and

$\mathscr{R}_e$ for the right charged-lepton closure singlets.

A charged-lepton closure map has the form

$\Gamma_e : \mathscr{R}_e \to \mathscr{L}_e$.

Written in a chosen basis, Γ_e is a matrix. But the matrix entries are not invariant: a different basis for $\mathscr{L}_e$ and $\mathscr{R}_e$ changes every entry.

Therefore, a raw matrix cannot be the first physical object.

1.2 The basic failure mode

Suppose a later paper writes

$$\Gamma = \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & k \end{pmatrix}$$

and then claims:

- a predicts the electron;
- e predicts the muon;
- k predicts the tau;
- b = 0 is a physical texture zero;
- the diagonal basis proves no mixing;
- a chosen rotation matrix is CKM.

All of these claims are invalid unless the basis has been fixed by a physical principle.

Under a frame rotation,

$\Gamma \mapsto U_L^\dagger \Gamma U_R$,

the entries change. A zero can become non-zero. A diagonal matrix can become non-diagonal. A phase can move. An eigenvector can rotate. The physics cannot be the raw coordinate table.

1.3 What survives

The following survive admissibly:

- singular values of Γ;
- eigenvalues of $\Gamma \Gamma^\dagger$;
- eigenvalues of $\Gamma^\dagger \Gamma$;
- the positive polar factor of Γ, as an operator class under conjugation;
- traces, determinants, and symmetric polynomials of Hermitian lifts;
- relative eigenframe data between two Hermitian lifts on a shared left frame;
- commutator invariants;
- spectral projectors;
- rephasing-invariant CP quantities;
- ratios of invariant eigenvalues;
- degeneracy-aware subspace relations.

That is the content of the Hermitian-lift firewall.

---

2. Named Premises of the Hermitian-Lift Theorem

HLF-1 — Occupancy-domain premise

Every closure map must have a stated Ω-address domain and codomain. A map without an Ω-address target is not an admissible VERSF prediction.

Falsifier: a mass or mixing claim targets "the electron," "the top," or "the quark sector" without specifying the occupied slot bundle and closure channel.

HLF-2 — Positive metric premise

A physical occupied bundle carries a positive inner product inherited from physical-sector positivity. In canonical kinetic coordinates this metric is the identity. If kinetic metrics are non-trivial but positive, the Hermitian lift must use the metric adjoint.

Falsifier: a Hermitian claim is made before a positive physical metric or canonical kinetic form is specified.

HLF-3 — Closure-map premise

A charged closure channel is represented by a linear map

$\Gamma_f : \mathscr{R}_f \to \mathscr{L}_f$

between right-terminal and left-attached occupied bundles. The map may be a Yukawa map, a mass map after electroweak symmetry breaking, or a VERSF closure-response map later shown to reduce to one of these.

Falsifier: a numerical closure claim is made without defining the map type.

HLF-4 — Frame-rotation covariance

Unitary frame changes act as

$\Gamma_f \mapsto U_L^\dagger \Gamma_f U_R$.

The right rotations U_R may be chosen independently per closure sector, since each right-terminal bundle is a separate Ω-address. The left rotations U_L may be chosen independently only for left bundles that are not weak partners within a single occupied doublet address; doublet-locked left bundles are governed by HLF-4′. The physics may not depend on the arbitrary choice of any admissible U_L or U_R.

Falsifier: a prediction changes under an admissible change of weak/flavour frame.

HLF-4′ — Doublet-locked frame premise

Admissible frame rotations must respect the Ω-address structure of SC-2. Left bundles that are weak components of a single occupied doublet address — $\mathscr{L}_u$ and $\mathscr{L}_d$ within Q_L; $\mathscr{L}_e$ and $\mathscr{L}_\nu$ within L_L — carry a single shared generation frame and admit only a common left rotation W. Independent left rotations of weak-partner bundles are not admissible frame changes.

This premise is load-bearing for mixing. If independent left rotations of weak partners were admissible, the up-sector and down-sector lifts could be diagonalised simultaneously by separate frame choices, and the relative frame V would dissolve into convention, failing the invariance test of §15.3 by this paper's own firewall logic. CKM is physical precisely because the doublet locks the two left frames together.

Falsifier: a derivation applies independent left generation rotations to the two weak components of one occupied doublet.

HLF-5 — Minimal Hermitian lift premise

The first admissible positive spectral object associated with Γ_f is quadratic and Hermitian:

$K_f^L = \Gamma_f \Gamma_f^\dagger$, $K_f^R = \Gamma_f^\dagger \Gamma_f$.

With non-trivial metrics, replace † by the metric adjoint ‡.

Falsifier: a raw non-Hermitian matrix spectrum is treated as a mass spectrum.

HLF-6 — Singular-spectrum premise

The non-zero eigenvalues of K_f^L and K_f^R are identical and equal to the squared singular values of Γ_f.

Falsifier: left and right lifts are claimed to produce different non-zero mass spectra for the same closure channel.

HLF-7 — Single-sector rotation firewall

A single closure sector has invariant spectral content but no absolute physical mixing frame.

Falsifier: the diagonalising matrix of a single isolated sector is treated as an observable mixing matrix.

HLF-8 — Shared-left-frame premise

A relative mixing matrix becomes meaningful only when two closure sectors act on a shared left weak-generation frame. For quarks, the up-type and down-type closures share the Q_L generation frame. For leptons, charged-lepton and neutrino closures share the L_L generation frame, but the neutrino closure exists only after a neutrino-mass extension is declared.

Falsifier: CKM or PMNS is claimed without two relevant left Hermitian lifts on a shared weak-generation frame.

HLF-9 — Degeneracy firewall

If a Hermitian lift has degenerate eigenvalues, eigenvectors inside the degenerate subspace are not unique. Only projectors onto degenerate subspaces and invariant relations between those projectors are physical.

Falsifier: a mixing angle inside an exactly degenerate subspace is treated as physical.

HLF-10 — Rephasing and permutation firewall

Eigenvectors are defined only up to phase and ordering. Physical claims must be rephasing-invariant and must state any ordering convention.

Falsifier: an arbitrary sign, phase, or row/column ordering is treated as physical.

HLF-11 — Texture firewall

A matrix-entry zero or texture pattern is not physical unless protected by a symmetry, invariant constraint, rank condition, or Ω-address selection rule.

Falsifier: a texture zero is claimed as a prediction while removable by a frame rotation.

HLF-12 — Extension separation

No neutrino Hermitian lift exists in the minimal massless-neutrino ledger. A PMNS-type relative-frame object requires a declared neutrino-mass extension.

Falsifier: PMNS is derived from the minimal ledger without adding a neutrino mass map.

HLF-13 — Invariant-expressibility premise

Every structural constraint offered as physical — a texture, an alignment, a vanishing, a degeneracy, an ordering — must be re-expressible as the vanishing or equality of a functional that is invariant under the allowed frame rotations, rephasings, and permutations. A constraint that cannot be so re-expressed is a coordinate statement, not a prediction.

Falsifier: a structural claim is asserted as physical while no frame-invariant reformulation of it is exhibited.

---

3. Occupied Bundles, Frames, and Closure Maps

3.1 Occupied bundles

For each closure sector f, define

ℒ_f

as the left-attached occupied generation bundle, and

$\mathscr{R}$_f

as the right-terminal occupied generation bundle.

For three generations, these are three-dimensional complex vector spaces, with basis labels inherited from the SC-1/SC-2 generation ledger.

Examples:

$\mathscr{R}_u = \mathrm{span}(u_R, c_R, t_R)$,

$\mathscr{R}_d = \mathrm{span}(d_R, s_R, b_R)$,

$\mathscr{R}_e = \mathrm{span}(e_R, \mu_R, \tau_R)$.

The corresponding left bundles are the relevant weak components inside Q_L or L_L.

3.2 Closure maps

The up-type closure map is

$\Gamma_u : \mathscr{R}_u \to \mathscr{L}_u$,

with Higgs route H̃.

The down-type closure map is

$\Gamma_d : \mathscr{R}_d \to \mathscr{L}_d$,

with Higgs route H.

The charged-lepton closure map is

$\Gamma_e : \mathscr{R}_e \to \mathscr{L}_e$,

with Higgs route H.

In the minimal ledger there is no map

$\Gamma_\nu : \mathscr{R}_\nu \to \mathscr{L}_\nu$

because there is no minimal right-handed sterile neutrino slot. A neutrino map requires an extension.

3.3 Frames

A **frame** is an orthonormal coordinate basis for an occupied generation bundle.

A frame rotation is a unitary change of that basis:

$\psi_L \mapsto U_L \psi_L$, $\psi_R \mapsto U_R \psi_R$.

The same physical closure map is then represented by a different coordinate matrix:

$\Gamma \mapsto U_L^\dagger \Gamma U_R$.

Therefore the matrix representation of Γ is not itself physical.

3.4 Shared weak-generation frames

The up and down left bundles are weak components of the same Q_L generation structure. Before closure diagonalisation, they inherit a common weak-generation frame.

Similarly, the charged-lepton and neutrino left bundles are weak components of the same L_L generation structure. PMNS becomes meaningful only if the neutrino side has its own mass/closure lift.

This is why CKM and PMNS are relative-frame objects, not single-sector objects.

The shared frame is not a descriptive convenience. It is enforced by HLF-4′, which types the common left rotation W as the only admissible left frame change on weak-partner bundles. Everything in §8 rests on this premise.

3.5 The shared generation factor

The statement that $\mathscr{L}_u$ and $\mathscr{L}_d$ "share a frame" must be made completely explicit, because §8 and §11 form commutators and projector overlaps between the two left lifts — and strictly, K_u^Q acts on the up weak component while K_d^Q acts on the down weak component. Operators on unrelated spaces have no commutator. The following lemma supplies the common space.

Shared Generation-Factor Lemma. For a weak doublet address, the occupied left bundle factorises as

$Q_L \cong \mathcal{G}_Q \otimes \mathbb{C}^2_{weak}$,

where $\mathcal{G}_Q$ is the shared generation factor, of dimension three, and $\mathbb{C}^2_{weak}$ carries the weak-component index with basis e_+, e_-. The up and down left bundles are the weak-component slices

$\mathscr{L}\_u = \mathcal{G}\_Q \otimes e_+$, $\mathscr{L}\_d = \mathcal{G}\_Q \otimes e_-$,

each canonically identified with $\mathcal{G}\_Q$, and hence with each other, through the weak-component slice maps. The left Hermitian lifts K_u^Q and K_d^Q are transported through this identification to endomorphisms of the single space $\mathcal{G}\_Q$. The commutator $[K\_u^Q, K\_d^Q]$, the projector overlaps $Tr(P\_a^u P\_b^d)$, and the CKM relative frame are defined on this common generation factor, not between unrelated coordinate spaces.

Proof. By I5, each occupied doublet slot carries an Ω-address containing a generation layer $i \in \{1, 2, 3\}$ and a weak-component label. By I6, the three generation copies are representation-identical, so the layer coordinate varies over identical representation data, while the weak-component label is fixed occupancy data. By I8, the doublet address therefore factorises as generation layer × weak component, and the occupied bundle is spanned by the independent (layer, component) pairs — the stated factorisation. The weak-component label is Ω-address data, fixed by occupancy, not a frame choice. Therefore the slice maps $\mathcal{G}\_Q \to \mathcal{G}\_Q \otimes e_+$ and $\mathcal{G}\_Q \to \mathcal{G}\_Q \otimes e_-$ are canonical, and composing one with the inverse of the other gives a canonical identification $\mathscr{L}\_u \cong \mathscr{L}\_d$ that no admissible frame rotation can alter. Under HLF-4′, admissible left rotations act as $W \otimes \mathbb{1}\_weak$ on the factorised bundle, so the identification is equivariant: both transported lifts conjugate by the same W on $\mathcal{G}\_Q$. ■

The same factorisation holds for L_L with generation factor $\mathcal{G}\_L$; the leptonic identification becomes load-bearing only when an extension supplies K_ν^L.

---

4. The Minimal Hermitian Lift and the Polar Decomposition

4.1 Canonical kinetic form

In canonical kinetic coordinates, the adjoint of Γ is the conjugate transpose Γ†.

The left Hermitian lift is

$K\_L = \Gamma \Gamma^\dagger : \mathscr{L} \to \mathscr{L}$.

The right Hermitian lift is

$K\_R = \Gamma^\dagger \Gamma : \mathscr{R} \to \mathscr{R}$.

Both are Hermitian:

$K\_L^\dagger = K\_L$, $K\_R^\dagger = K\_R$.

Both are positive semidefinite:

$\langle x, K_L x\rangle = \langle \Gamma^\dagger x, \Gamma^\dagger x\rangle \geq 0$,

$\langle y, K_R y\rangle = \langle \Gamma y, \Gamma y\rangle \geq 0$.

4.2 Metric-adjoint form

If the occupied bundles carry non-trivial positive metrics G_L and G_R, the correct adjoint is the metric adjoint ‡, defined by

$\langle \Gamma r, l\rangle_L = \langle r, \Gamma^\ddagger l\rangle_R$.

In coordinates,

$\Gamma^\ddagger = G_R^{-1} \Gamma^\dagger G_L$.

Then

$K_L = \Gamma \Gamma^\ddagger$, $K_R = \Gamma^\ddagger \Gamma$.

In canonical coordinates, $G_L = G_R = \mathbb{1}$, and ‡ = †. Positivity and self-adjointness hold with respect to the physical metrics, not with respect to an arbitrary coordinate inner product; this is why HLF-2 is a load-bearing premise rather than a formality.

Canonicalisation Equivalence Lemma. Define the canonicalised map

$\hat{\Gamma} = G_L^{(1/2)} \Gamma G_R^{(-1/2)}$.

Then

$\hat{\Gamma} \hat{\Gamma}^\dagger = G_L^{(1/2)} (\Gamma \Gamma^\ddagger) G_L^{(-1/2)}$, $\hat{\Gamma}^\dagger \hat{\Gamma} = G_R^{(1/2)} (\Gamma^\ddagger \Gamma) G_R^{(-1/2)}$,

so the metric-adjoint lifts are similar to the ordinary Hermitian lifts of $\hat{\Gamma}$. Their spectra coincide, are real and non-negative, and the physical singular values are those of $\hat{\Gamma}$. Canonicalising the kinetic terms first and forming †-lifts, or keeping the metrics and forming ‡-lifts, are the same operation up to positive similarity.

Proof. Direct substitution: $\Gamma \Gamma^\ddagger = \Gamma G_R^{-1} \Gamma^\dagger G_L$, and $G_L^{(1/2)} (\Gamma \Gamma^\ddagger) G_L^{(-1/2)} = G_L^{(1/2)} \Gamma G_R^{-1} \Gamma^\dagger G_L^{(1/2)} = \hat{\Gamma} \hat{\Gamma}^\dagger$; symmetrically on the right. Similar operators share characteristic polynomials, and $\hat{\Gamma} \hat{\Gamma}^\dagger$ is Hermitian positive semidefinite in the canonical inner product, so the common spectrum is real and non-negative. ■

This closes a specific misreading in advance: no coordinate eigenvalue of a non-self-adjoint matrix representation is ever the spectral target. The target is the spectrum of the metric-self-adjoint lift, equivalently of the canonicalised Hermitian lift — the same numbers by the lemma.

4.3 Why Hermitian

Masses and Yukawa magnitudes are spectral observables.

A general complex matrix Γ is not Hermitian and need not have real eigenvalues. Its eigenvalues are not stable mass targets: they are not invariant under the independent left/right frame rotations of HLF-4, and for a non-normal matrix they do not even coincide with its singular values. Its entries are basis-dependent. Its diagonal form may be a convention.

The Hermitian lift repairs this. It produces a positive operator whose eigenvalues are real, non-negative, and invariant under frame rotations.

4.4 The polar decomposition

The Hermitian lift is not an ad hoc repair. It is the exact frame-invariant "magnitude" part of the closure map.

Lemma 0 — Polar factorisation. Every closure map $\Gamma : \mathscr{R} \to \mathscr{L}$ admits a polar decomposition

$\Gamma = P\, W$,

where $P = \sqrt{(\Gamma \Gamma^\dagger)}$ is positive semidefinite on $\mathscr{L}$ and $W : \mathscr{R} \to \mathscr{L}$ is the canonical partial isometry with initial space $(\ker \Gamma)^\perp \subseteq \mathscr{R}$ and final space $\mathrm{ran}(P) \subseteq \mathscr{L}$. Both factors are unique: P by positivity of the square root, and W because it is determined by Γ on $(\ker \Gamma)^\perp$ and vanishes on $\ker(\Gamma) \subseteq \mathscr{R}$. When $\mathscr{R}$ and $\mathscr{L}$ have equal dimension and Γ is invertible, W is a unique unitary. For rectangular channels, the non-zero spectra of K_L and K_R still coincide, while their zero-eigenvalue multiplicities differ by the dimension mismatch.

Proof. $\Gamma \Gamma^\dagger$ is positive semidefinite by §4.1, so its unique positive square root P exists. Define $W = P^{-1} \Gamma$ on $(\ker \Gamma)^\perp = \mathrm{ran}(\Gamma^\dagger)$, with P^{-1} the inverse on $\mathrm{ran}(P)$, and $W = 0$ on $\ker(\Gamma)$; then $P\, W = \Gamma$, and $W^\dagger W$ is the orthogonal projector onto $(\ker \Gamma)^\perp$, so W is the canonical partial isometry, unitary in the square invertible case. Uniqueness of P follows from uniqueness of the positive square root: if $\Gamma = P' W'$ with $P' \geq 0$ and W' a partial isometry with the stated initial space, then $\Gamma \Gamma^\dagger = P'^2$, so $P' = P$, and $W' = W$ follows on $(\ker \Gamma)^\perp$ from $P\, W' = \Gamma$ and on $\ker(\Gamma)$ from the initial-space convention. The coincidence of non-zero spectra of K_L and K_R holds for any rectangular Γ, since $\Gamma^\dagger \Gamma$ and $\Gamma \Gamma^\dagger$ share all non-zero eigenvalues. ■

Interpretation. The polar decomposition splits Γ into an invariant-magnitude part P and a pure-frame part W. Under $\Gamma \mapsto U_L^\dagger \Gamma U_R$, the positive factor transforms by conjugation, $P \mapsto U_L^\dagger P U_L$, while the partial-isometry factor absorbs the entire arbitrary frame content, $W \mapsto U_L^\dagger W U_R$. The Hermitian lift $K_L = P^2$ therefore captures exactly what survives when the frame freedom is quotiented away. Nothing physical lives in W alone.

All minimal SF-1 channels are square three-by-three maps, so the unitary form suffices inside this gate. Extension gates that introduce rectangular closure blocks — seesaw sectors, vector-like additions — must use the partial-isometry form and must not assume equal zero-eigenvalue multiplicities for the two lifts.

4.5 Why minimal

The lift is minimal in four senses.

First, it is the lowest-degree non-trivial positive object built from Γ and its adjoint: no expression linear in Γ is simultaneously an endomorphism of a single occupied bundle, positive, and frame-covariant.

Second, it acts on the same occupied left or right bundle, rather than enlarging the theory.

Third, it preserves the singular spectrum of Γ without adding arbitrary structure.

Fourth, it is forced by covariance: to eliminate the arbitrary right frame from a left operator, the right index must be contracted with its adjoint, producing $\Gamma \Gamma^\dagger$. To eliminate the arbitrary left frame from a right operator, the left index must be contracted, producing $\Gamma^\dagger \Gamma$. Up to positive normalisation, no other quadratic contraction exists.

Block-Hermitian auxiliary firewall. One linear Hermitian construction does exist, on a doubled space:

$$\mathcal{D}\_\Gamma = \begin{pmatrix} 0 & \Gamma \\ \Gamma^\ddagger & 0 \end{pmatrix},$$

acting on $\mathscr{L} \oplus \mathscr{R}$, with spectrum $\pm\sigma\_i$. This does not compete with the minimal lift, for three reasons. It is not positive. It does not act on a single occupied Ω-address bundle, but on an auxiliary doubled space that no occupancy assignment supplies. And its square is block diagonal,

$$\mathcal{D}\_\Gamma^2 = \begin{pmatrix} \Gamma \Gamma^\ddagger & 0 \\ 0 & \Gamma^\ddagger \Gamma \end{pmatrix},$$

so its invariant positive content is exactly the pair of Hermitian lifts already identified. $\mathcal{D}\_\Gamma$ is an admissible auxiliary representation — occasionally convenient for computation — but not a separate primitive prediction target. Its negative spectral branch is a bookkeeping artifact of the doubling, never a physical mass.

Scope of the uniqueness claim. Minimality and uniqueness are asserted relative to the declared structure: the Ω-address bundles, the physical metrics, and nothing else. If a later gate supplies an additional symmetry-fixed tensor — a projector, a charge operator, a selection tensor, a flavour-breaking spurion — then further covariant quadratic expressions become available, and their admissibility is governed by HLF-13, not overruled by this clause. Absent such declared extra structure, the minimal positive lift is uniquely $\Gamma \Gamma^\ddagger$ or $\Gamma^\ddagger \Gamma$.

4.6 Relation to singular values

If

$\Gamma = U\_L \Sigma U\_R^\dagger$

is a singular-value decomposition, then

$K\_L = U\_L \Sigma^2 U\_L^\dagger$, $K\_R = U\_R \Sigma^2 U\_R^\dagger$.

The non-zero eigenvalues of both lifts are

σ_1^2 , σ_2^2 , σ_3^2 .

The singular values σ_i are the admissible closure magnitudes.

After electroweak symmetry breaking, if Γ_f is a Yukawa map Y_f, then

$m\_i = (v / \sqrt{2})\, \sigma\_i(Y\_f)$.

SF-1 does not derive v or σ_i. It defines the admissible target.

---

5. Frame Rotations and Covariance

5.1 Transformation of the closure map

Under frame rotations,

$\Gamma \mapsto \Gamma' = U\_L^\dagger \Gamma U\_R$.

Then

$K\_L' = \Gamma' \Gamma'^\dagger = U\_L^\dagger \Gamma U\_R U\_R^\dagger \Gamma^\dagger U\_L = U\_L^\dagger K\_L U\_L$,

and similarly

$K\_R' = \Gamma'^\dagger \Gamma' = U\_R^\dagger K\_R U\_R$.

Thus each Hermitian lift transforms by unitary conjugation.

5.2 What conjugation preserves

Unitary conjugation preserves:

- eigenvalues;
- rank;
- trace;
- determinant;
- characteristic polynomial;
- degeneracy structure;
- spectral projectors, as covariant operator classes;
- commutator traces between commonly transformed operators.

It does not preserve:

- individual matrix entries;
- whether a matrix looks diagonal;
- arbitrary texture zeros;
- eigenvector coordinates;
- phase conventions;
- row and column labels, unless fixed by an ordering convention.

5.3 Lemma 1 — Spectral covariance

Lemma 1. The spectra of K_L and K_R are invariant under arbitrary unitary frame rotations of the left and right occupied bundles.

Proof. Under frame rotation, the lifts transform by unitary conjugation:

$$K_L' = U_L^\dagger K_L U_L , \quad K_R' = U_R^\dagger K_R U_R .$$

Unitary conjugation preserves the characteristic polynomial and hence the eigenvalues. Therefore the spectra are frame-invariant. ■

5.4 Lemma 2 — Raw matrix entries are not invariant

Lemma 2. A generic entry Γ_{ij} of a closure matrix is not a frame-invariant quantity.

Proof. Under $\Gamma \mapsto U_L^\dagger \Gamma U_R$, the transformed entry is

$$\Gamma'_{ij} = \Sigma_{\{a,b\}} (U_L^\dagger)_{ia} \Gamma_{ab} (U_R)_{bj} .$$

For generic U_L, U_R this is a linear combination of many entries of the original matrix. More strongly: $\Gamma'_{ij} = x^\dagger \Gamma y$ with $x = U_L e_i$ and $y = U_R e_j$, and as the frames vary, x and y range independently over all unit vectors. The attainable set at any position (i, j) is therefore exactly the disk $|z| \leq \sigma_{max}(\Gamma)$ — the maximum is the largest singular value, zero is attained by orthogonality on bundles of dimension at least two (all occupied generation bundles are three-dimensional), and intermediate moduli follow by continuity, with phases free throughout. Every non-zero closure map thus admits frames realising a continuum of values at any chosen position, and Γ_ij has no frame-independent meaning unless additional structure fixes the frames or protects the entry. ■

---

6. The Raw-Matrix Firewall and Invariant Expressibility

6.1 Forbidden prediction targets

The following are not admissible physical prediction targets by themselves:

Γ_{12} , Γ_{23} , Γ_{33} ,

or any other raw matrix entry.

Likewise, the following are not admissible unless symmetry-fixed:

$\Gamma_{12} = 0$, $\Gamma_{13} = \Gamma_{31}$, Γ diagonal , Γ triangular , Γ of a chosen texture .

These are coordinate statements.

6.2 The invariant-expressibility rule

HLF-13 upgrades the texture firewall from a prohibition into a constructive test. A structural claim becomes admissible exactly when it can be rewritten as an invariant condition. Admissible reformulations include:

1. a symmetry that fixes the frame, with the residual frame group stated;
2. an Ω-address selection rule that forbids a closure route;
3. a rank constraint, e.g. $\mathrm{rank}(\Gamma) \leq 2$, equivalently $\det(K) = 0$;
4. an invariant vanishing condition, e.g. $\det(\Gamma) = 0$ or an eigenvalue equality;
5. a projector relation, e.g. $\mathrm{Tr}(K P) = 0$ for a symmetry-fixed projector P;
6. a commutator relation between commonly framed lifts;
7. a conserved quantum number that blocks a coupling.

The test cuts both ways. A rank-one claim is admissible because rank is invariant. An isolated statement such as "the 1–3 entry is zero" is not — unless the claimant exhibits the invariant functional whose vanishing it expresses.

Throughout this rule, "invariant under the allowed frame rotations" means invariance under the doublet-locked group of HLF-4 and HLF-4′: independent right rotations, independent left rotations of unlocked bundles, and a single common left rotation on weak-partner bundles of one doublet address.

6.3 Diagonalisation is not prediction

Every Hermitian matrix can be diagonalised. Therefore, writing

$K = U D U^\dagger$

does not by itself predict anything beyond the spectrum D.

Similarly, every single Yukawa matrix can be singular-value decomposed. That does not create CKM. CKM requires comparing two left diagonalisation frames attached to the same weak-generation structure.

6.4 Lemma 3 — Texture-zero firewall

Lemma 3. A matrix-entry zero is not physically meaningful unless preserved by the allowed frame rotations or re-expressible as a frame-invariant condition.

Proof. By Lemma 2, a generic matrix entry changes under unitary frame rotations. A zero entry in one frame can become non-zero in another. Therefore a texture zero is not invariant under the allowed frame freedom. It becomes physical only if the allowed frame freedom is reduced by symmetry, selection rule, or ordering convention, or if the zero is equivalent to the vanishing of an invariant functional as in §6.2. ■

---

7. Spectral Content and Mass/Yukawa Targets

7.1 Valid spectral targets

A later VERSF mass or Yukawa theorem may target

$\mathrm{spec}(K_f)$,

or invariant functions such as

$\mathrm{Tr}(K_f)$, $\det(K_f)$, $\lambda_1(K_f)/\lambda_2(K_f)$, $\sqrt{\lambda_i}(K_f)/\sqrt{\lambda_j}(K_f)$.

It may also target hierarchy exponents if they are explicitly tied to eigenvalue ratios:

$\lambda_i(K_f) \sim \varepsilon^{n_i}$,

with the ordering convention for i stated per HLF-10.

Any such gate must first type ε under QN-1 discipline: ε must be declared as a frame-invariant, Ω-addressed quantity — an eigenvalue ratio, an invariant coupling, or another Hermitian-lift observable — before exponents are assigned to it. An untyped ε is exactly the free fitting knob the κ-ledger was built to exclude.

It may not target a raw entry unless the frame has been physically fixed.

7.2 Masses from the lift

If $\Gamma_f = Y_f$ is a Yukawa map, then

$K_f = Y_f Y_f^\dagger$

has eigenvalues

y_1^2 , y_2^2 , y_3^2 .

After electroweak symmetry breaking,

$m_i^2 = (v^2/2)\, y_i^2$.

Therefore, masses correspond to square roots of Hermitian-lift eigenvalues, not eigenvalues of an arbitrary non-Hermitian Yukawa matrix.

7.3 Lemma 4 — Single-sector physical content

Lemma 4. For a single isolated closure sector, the frame-invariant content is exhausted by the singular values of Γ_f, equivalently the eigenvalues of its Hermitian lifts.

Proof. By singular-value decomposition, $\Gamma_f = U_L \Sigma U_R^\dagger$ with Σ diagonal and non-negative. The admissible rotations of HLF-4 and HLF-4′ still act transitively on the unitary factors of a single sector — the common left rotation W is unrestricted when only one sector's map is in play, since the lock constrains joint transformations of weak partners, not the rotation itself: for any target unitaries V_L, V_R, the frame choice $U_L' = U_L V_L^\dagger$, $U_R' = U_R V_R^\dagger$ replaces (U_L, U_R) by (V_L, V_R) while leaving Σ fixed. Hence every function of Γ_f that is invariant under all allowed frame rotations is a function of Σ alone. Conversely Σ is invariant by Lemma 1. Therefore a single isolated closure sector carries exactly its singular spectrum and no absolute mixing-frame content. ■

---

8. Relative Frames, CKM, and the Parameter Count

8.1 Two lifts on the quark weak-generation frame

The up and down sectors have closure maps

$\Gamma_u : \mathscr{R}_u \to \mathscr{L}_u$, $\Gamma_d : \mathscr{R}_d \to \mathscr{L}_d$.

Their left spaces $\mathscr{L}_u$ and $\mathscr{L}_d$ are the upper and lower weak components of the same quark doublet generation structure Q_L. By the Shared Generation-Factor Lemma of §3.5, both are canonically identified with the common generation factor $\mathcal{G}_Q$, so the two left Hermitian lifts are transported to endomorphisms of one space:

$K_u^Q = \Gamma_u \Gamma_u^\dagger$, $K_d^Q = \Gamma_d \Gamma_d^\dagger$, both acting on $\mathcal{G}_Q$.

Every comparison in this section — diagonalisation mismatch, projector overlap, commutator — is performed on $\mathcal{G}_Q$. Without the identification, the objects below would be ill-typed.

8.2 Diagonalisation

Let

$U_{uL}^\dagger K_u^Q U_{uL} = D_u^2$, $U_{dL}^\dagger K_d^Q U_{dL} = D_d^2$.

The relative rotation is

$V_{CKM} = U_{uL}^\dagger U_{dL}$.

This is not the absolute orientation of U_{uL} or U_{dL}. It is the mismatch between them.

8.3 Why relative rotation is physical

By HLF-4′, the only admissible left frame change of the pair is a common quark weak-frame rotation W_Q, under which both left lifts transform by conjugation:

$K_u^Q \mapsto W_Q^\dagger K_u^Q W_Q$, $K_d^Q \mapsto W_Q^\dagger K_d^Q W_Q$.

Their diagonalising frames transform together: $U_{uL} \mapsto W_Q^\dagger U_{uL}$ and $U_{dL} \mapsto W_Q^\dagger U_{dL}$, so the product $U_{uL}^\dagger U_{dL}$ is unchanged. The relative rotation survives common frame rotation exactly, and survives the residual per-sector freedom up to rephasing and ordering conventions.

Thus CKM is frame-relative, not frame-absolute — and it is HLF-4′ that makes the relative object well defined. Were independent left rotations of $\mathscr{L}_u$ and $\mathscr{L}_d$ admissible, both lifts could be diagonalised at once and V would be pure convention.

8.4 Residual freedom and the parameter count

Lemma 5 — Residual freedom of the relative frame. Suppose K_u^Q and K_d^Q have non-degenerate spectra and an eigenvalue-ordering convention is fixed. Then the diagonalising unitaries are determined up to diagonal phase matrices, $U_{uL} \mapsto U_{uL} \Phi_u$ and $U_{dL} \mapsto U_{dL} \Phi_d$, and the relative frame transforms as

$$V \mapsto \Phi_u^\dagger V \Phi_d .$$

The physical content of V is therefore the rephasing class of V.

Proof. With non-degenerate spectrum and fixed ordering, each eigenvector is determined up to a phase, so the residual per-sector frame group is the diagonal torus. The transformation of V follows by substitution. ■

Lemma 5′ — Relative-frame parameter count. For n non-degenerate generations, with full diagonal rephasing tori on both sectors — the Dirac-type case — the rephasing class of an $n \times n$ unitary relative-frame matrix carries

$$n(n-1)/2 \text{ angles and } (n-1)(n-2)/2 \text{ phases.}$$

For $n = 3$ this is three angles and one phase. For $n = 2$ it is one angle and no phase. Majorana-type extensions reduce one torus to signs and enlarge the phase count, as recorded in §9.3.

Proof. A unitary $n \times n$ matrix has n^2 real parameters. The residual rephasing group $\Phi_u^\dagger (\cdot) \Phi_d$ has 2n phases, of which one overall common phase acts trivially, removing $2n - 1$ parameters. This leaves $n^2 - (2n - 1) = (n - 1)^2$ physical parameters. Of these, $n(n-1)/2$ are angles of an orthogonal rotation, and the remainder, $(n-1)^2 - n(n-1)/2 = (n-1)(n-2)/2$, are phases. ■

Consequence. Any later VERSF CKM theorem for three generations must deliver exactly four invariant quantities — three angles and one CP-sensitive phase — or an equivalent invariant set such as four independent $|V_{ab}|^2$ values consistent with unitarity, or three $|V_{ab}|^2$ values plus one quartet invariant. A claim delivering more independent numbers than the invariant count has smuggled in a frame convention.

8.5 Projector formulation

Let P_a^u be the spectral projector onto the a-th up-sector eigenstate, and P_b^d the projector onto the b-th down-sector eigenstate. Then in the non-degenerate case,

$$\mathrm{Tr}(P_a^u P_b^d) = |V_{ab}|^2 .$$

This expression is manifestly invariant under common frame rotation and under all residual rephasings.

Therefore, the physical content of CKM magnitudes is the relative geometry of spectral projectors.

8.6 Lemma 6 — CKM requires two left lifts

Lemma 6. A CKM-type mixing matrix cannot be obtained from a single closure sector. It requires two Hermitian lifts acting on a shared left weak-generation frame.

Proof. By Lemma 4, a single sector's invariant content is exhausted by its singular spectrum; its diagonalising frame is pure convention. With two lifts on a shared left frame, HLF-4′ forbids the independent left rotations that would diagonalise both simultaneously, so no admissible frame choice generally diagonalises the pair. The relative mismatch between their eigenframes is invariant under the admissible common rotation by §8.3, and its rephasing class is invariant under all residual freedom by Lemma 5. This rephasing class is the CKM-type object. ■

8.7 Corollary — No mixing if the lifts commute

If

$[K_u^Q, K_d^Q] = 0$

and the relevant spectra are non-degenerate, then K_u^Q and K_d^Q are simultaneously diagonalisable. The relative mixing is trivial up to rephasing and permutation.

Therefore, non-trivial CKM structure requires non-commuting left Hermitian lifts. The converse discipline also holds: a later paper may not simultaneously claim a non-zero commutator invariant and a trivial mixing matrix, or a vanishing commutator with non-degenerate spectra and a non-trivial mixing matrix.

---

9. Lepton Frames and PMNS Extension Dependence

9.1 Charged-lepton lift

The charged-lepton closure map is

$\Gamma_e : \mathscr{R}_e \to \mathscr{L}_e$.

Its left Hermitian lift is

$K_e^L = \Gamma_e \Gamma_e^\dagger$.

This supplies the charged-lepton spectral frame.

9.2 Minimal neutrino ledger

In the minimal chiral matter ledger, the left-handed neutrino component exists:

$\nu_L \subset L_L$.

But no minimal right-handed sterile neutrino slot exists. Therefore no minimal Dirac neutrino closure map is present. Nor has a Majorana or Weinberg-operator closure been declared inside SF-1.

Thus there is no minimal neutrino Hermitian lift in SF-1.

9.3 PMNS requires an extension

A PMNS-type object requires a neutrino lift

K_ν^L .

This may arise from an extension such as:

- Dirac neutrino closure;
- Majorana mass closure;
- seesaw extension;
- Weinberg operator;
- a VERSF-specific neutrino completion gate.

Once such a lift exists, one may define

$U_{eL}^\dagger K_e^L U_{eL} = D_e^2$, $U_{\nu L}^\dagger K_\nu^L U_{\nu L} = D_\nu^2$,

and

$U_{PMNS} = U_{eL}^\dagger U_{\nu L}$.

Without K_ν^L, there is no PMNS target in the minimal ledger.

One extension-sensitive refinement is recorded now, so that no later paper can overlook it: if the declared extension makes the neutrino mass operator Majorana, the residual rephasing freedom on the neutrino side is reduced from a full diagonal torus to diagonal signs. The rephasing class of U_PMNS is then larger than the Dirac-case class, and carries the additional invariant phases usually called Majorana phases. SF-1 does not derive them; it notes that the invariant class of a PMNS claim depends on the declared extension type, and that this dependence must be stated in the extension ledger.

A second refinement runs deeper than the rephasing class. A Majorana mass operator is a symmetric bilinear on $\mathscr{L}_\nu$ — $m_\nu = m_\nu^T$ in every frame — so there is no independent right bundle at all. The relevant factorisation is Takagi,

$$m_\nu = U D U^T ,$$

with D diagonal and non-negative, not a singular-value decomposition, and the frame freedom is a single unitary acting on one side. The left Hermitian lift remains well defined, $K_\nu^L = m_\nu m_\nu^\dagger$, and its eigenvalues remain the admissible spectral targets. But the "right lift" of the Dirac machinery is structurally absent, not merely extension-dependent. A later extension gate that imports the SVD left/right apparatus wholesale into a Majorana channel produces a malformed lift and fails this gate.

The invariant class is then countable per extension type. For a Dirac three-generation lepton extension, Lemma 5′ applies unchanged: PMNS carries three angles and one Dirac-type phase, the same invariant class as CKM. For a Majorana extension, the neutrino-side rephasing torus is reduced to signs, so only the charged-lepton torus is quotiented; the class then carries three angles and three phases — one Dirac-type phase and two Majorana phases. SF-1 does not derive any of these values; it fixes the admissible invariant class that later extension papers must respect, and F24 fires on any claim that uses the wrong class for its declared extension.

9.4 Lemma 7 — PMNS extension dependence

Lemma 7. PMNS-type mixing is not available in the minimal massless-neutrino occupancy ledger. It becomes an admissible relative-frame target only after a neutrino-mass extension supplies a neutrino Hermitian lift on the shared lepton weak-generation frame, and its invariant class is fixed by the extension type.

Proof. By SC-2, ν_L exists as the upper component of L_L, but ν_R is absent from the minimal census. Without a neutrino closure map or effective Majorana operator, there is no K_ν^L. A relative-frame matrix requires two left lifts on the shared L_L frame; the charged-lepton lift alone supplies spectral content but no lepton mixing matrix by Lemma 4. When an extension supplies K_ν^L, Lemma 5 applies with the residual neutrino-side freedom determined by the operator type — full torus for Dirac, signs for Majorana — so the invariant class of the relative frame is extension-dependent. ■

---

10. Degeneracy, Rephasing, and Permutation Firewalls

10.1 Degenerate eigenvalues

If a Hermitian lift has distinct eigenvalues, its eigenvectors are fixed up to phases and ordering.

If two eigenvalues are exactly equal, any unitary rotation inside the degenerate subspace leaves the lift unchanged. Therefore, eigenvectors inside a degenerate subspace are not physical. The physical object is the projector onto the degenerate subspace.

10.2 Degeneracy-safe prediction

A valid VERSF prediction may say

$\lambda_1 = \lambda_2$

or

$\mathrm{rank}(K) = 1$.

These are invariant.

But it may not claim a definite mixing angle inside an exactly degenerate subspace unless additional structure breaks the degeneracy. In the presence of degeneracy, the projector formulation of §8.5 remains valid at the level of subspace projectors: $\mathrm{Tr}(P_S^u P_T^d)$ is invariant for degenerate subspaces S and T, and is the correct degeneracy-safe replacement for $|V_{ab}|^2$.

10.3 Rephasing

Eigenvectors may be multiplied by arbitrary phases:

$|u_i\rangle \mapsto e^{i\alpha_i} |u_i\rangle$.

Mixing matrices therefore change by row and column phases. Physical claims must be rephasing-invariant.

Examples include

$|V_{ij}|$,

quartet invariants such as

$Q_{(ij,kl)} = V_{ij} V_{kl} V_{il} V_{kj}$,

and commutator invariants.

10.4 Permutation

Eigenvalue ordering is a convention until tied to an invariant ordering rule, such as ascending eigenvalue, descending eigenvalue, or a later hierarchy theorem. A prediction must state which ordering convention is used.

This is also the discipline that blocks generation-label drift: "generation 1" is not an invariant name unless it is tied to a stated position in the ordered Hermitian-lift spectrum.

10.5 Lemma 8 — Degenerate-frame firewall

Lemma 8. Inside an exactly degenerate eigenspace of a Hermitian lift, no individual eigenvector direction is physical without additional structure.

Proof. Let K have eigenvalue λ on a subspace S of dimension greater than one. For any unitary U_S acting inside S,

K restricted to S equals $\lambda \mathbb{1}_S$

and is unchanged by U_S. Therefore the choice of eigenvectors inside S is arbitrary. Only the projector onto S is invariant. ■

---

11. CP, Commutator Invariants, and the Quartet Identity

11.1 CP cannot be read from one arbitrary matrix phase

A complex phase in a raw matrix entry is not automatically a CP prediction. Phases can move under rephasing and frame rotation. A CP claim must target an invariant.

11.2 Quark-sector commutator

For the quark sector, define

$C_{ud} = [K_u^Q, K_d^Q]$,

well typed as an operator on the shared generation factor $\mathcal{G}_Q$ per §3.5. Under common quark-frame rotation,

$$C_{ud} \mapsto W_Q^\dagger C_{ud} W_Q .$$

Therefore traces of powers of C_{ud}, its determinant, and related scalar functions are frame-invariant. Note also that C_{ud} is anti-Hermitian, so $\mathrm{Tr}(C_{ud}) = 0$ identically and $\det(C_{ud})$ is purely imaginary in odd dimension. For a traceless three-by-three matrix, $\det(C_{ud}) = \mathrm{Tr}(C_{ud}^3)/3$, so the cubic trace and the determinant are the same invariant. The independent low-degree scalar invariants are therefore $\mathrm{Tr}(C_{ud}^2)$, which is CP-even, and $\det(C_{ud})$, which carries the CP-odd content.

11.3 Proposition — the commutator–quartet identity

Proposition 1. Let K_u^Q and K_d^Q be three-generation left Hermitian lifts with eigenvalues λ_a^u and λ_b^d, and let V be their relative-frame matrix. Then

$$\det(C_{ud}) = 2i\, J \cdot \Pi_{(a<a')} (\lambda_{a'}^u - \lambda_a^u) \cdot \Pi_{(b<b')} (\lambda_{b'}^d - \lambda_b^d) ,$$

where

$$J = \mathrm{Im}(V_{11}\, V_{22}\, V_{12}\, V_{21})$$

is the rephasing-invariant quartet, with any single quartet $Q_{(ij,kl)}$ determining the same $|J|$ by unitarity. Both sides are invariant under common frame rotation and all residual rephasing.

Proof. Work in the frame that diagonalises K_u^Q, admissible as a common left rotation per HLF-4′, and expand the determinant of the commutator. The full derivation is carried out inside the corpus in Appendix A §A.1, using only unitarity and a Vandermonde divisibility argument. The identity is structural: it constrains the invariant class of CP claims and asserts no numerical value. Sign conventions for J follow the stated ordering convention per HLF-10. ■

Consequences.

1. CP-sensitive content vanishes if and only if $J = 0$ or any two eigenvalues within one sector coincide.
2. A later VERSF CP theorem has exactly one independent CP-odd invariant to deliver in the three-generation quark sector.
3. Any claimed CP prediction must be checkable against $\det(C_{ud})$, which is computable from the two lifts without ever fixing a frame.

11.4 Lemma 9 — CP phase firewall

Lemma 9. A raw complex phase in a closure matrix is not an admissible CP prediction unless it appears in a rephasing-invariant, frame-invariant quantity.

Proof. Under independent left and right rephasings, phases of individual matrix entries change. Therefore an entry phase is not physical. CP-eligible content must survive the allowed frame and rephasing transformations. By Proposition 1, the surviving CP-odd content in the three-generation quark sector is carried by det(C_ud), equivalently by J together with the eigenvalue splittings. Commutator invariants and quartet invariants satisfy the requirement; raw phases do not. ■

---

12. Theorem 1 — The Minimal Hermitian Lift Theorem

12.1 Statement

Theorem 1 — Minimal Hermitian Lift Theorem. Let

$\Gamma_f : \mathscr{R}_f \to \mathscr{L}_f$

be an admissible VERSF closure map between occupied Ω-address bundles, with positive physical inner products on $\mathscr{R}_f$ and $\mathscr{L}_f$. Then the minimal admissible positive spectral objects associated with Γ_f are the Hermitian lifts

$K_f^L = \Gamma_f \Gamma_f^‡ , K_f^R = \Gamma_f^‡ \Gamma_f ,$

where ‡ is the metric adjoint. In canonical kinetic coordinates these reduce to

$K_f^L = \Gamma_f \Gamma_f^† , K_f^R = \Gamma_f^† \Gamma_f .$

The lifts are:

1. self-adjoint with respect to the physical metrics;
2. positive semidefinite;
3. covariant under frame rotations;
4. spectrum-invariant under unitary frame choice;
5. singular-spectrum equivalent, sharing all non-zero eigenvalues;
6. equal to the squared positive polar factors of Γ_f, hence carrying exactly the frame-invariant magnitude content of the closure map;
7. minimal and unique, up to positive normalisation and the left/right choice, among quadratic positive frame-covariant constructions from Γ_f — relative to the declared Ω-address bundles and physical metrics, and in the absence of additional symmetry-fixed tensors, spurions, or extension data;
8. admissible targets for later mass, Yukawa, and hierarchy predictions.

12.2 Proof

By HLF-1 and HLF-3, Γ_f maps a right-terminal occupied bundle into a left-attached occupied bundle. By HLF-2, these bundles carry positive physical inner products. Therefore the metric adjoint $\Gamma_f^\ddagger$ exists.

Define $K_f^L = \Gamma_f \Gamma_f^\ddagger$. For any $l \in \mathscr{L}_f$,

$$\langle l, K_f^L l\rangle_L = \langle l, \Gamma_f \Gamma_f^\ddagger l\rangle_L = \langle \Gamma_f^\ddagger l, \Gamma_f^\ddagger l\rangle_R \geq 0 .$$

Thus K_f^L is positive semidefinite, and it is self-adjoint by construction. The same argument applies to $K_f^R = \Gamma_f^\ddagger \Gamma_f$. In canonical coordinates $\ddagger = \dagger$, giving the familiar forms.

Under frame rotations, $\Gamma_f \mapsto U_L^\dagger \Gamma_f U_R$, and the lifts transform as

$$K_f^L \mapsto U_L^\dagger K_f^L U_L , \quad K_f^R \mapsto U_R^\dagger K_f^R U_R ,$$

so each lift is frame-covariant and its spectrum is invariant under frame choice by Lemma 1.

The non-zero spectra of K_f^L and K_f^R coincide because both are generated by the same singular values of Γ_f: in singular-value form $\Gamma_f = U_L \Sigma U_R^\dagger$, so

$$K_f^L = U_L \Sigma^2 U_L^\dagger , \quad K_f^R = U_R \Sigma^2 U_R^\dagger .$$

By Lemma 0, $K_f^L = P^2$ where P is the unique positive polar factor of Γ_f, establishing point 6.

For minimality and uniqueness, classify the candidate constructions. A construction linear in Γ_f alone is a map $\mathscr{R}_f \to \mathscr{L}_f$, not an endomorphism of either bundle, so it has no frame-stable spectrum. The one linear Hermitian construction available — the block operator $\mathcal{D}_\Gamma$ of §4.5 on the doubled space $\mathscr{L}_f \oplus \mathscr{R}_f$ — is not positive, does not act on an occupied Ω-address bundle, and squares to the block-diagonal pair of quadratic lifts; it is an auxiliary representation, not a competing primitive. A construction linear in Γ_f and $\Gamma_f^\ddagger$ jointly that produces a left endomorphism must contract the single free right index of Γ_f against the single free right index of $\Gamma_f^\ddagger$; the only such quadratic contraction is $\Gamma_f \Gamma_f^\ddagger$ up to positive scalar. Symmetrically, the only right-endomorphism contraction is $\Gamma_f^\ddagger \Gamma_f$. Higher-degree positive covariant constructions, such as $(\Gamma_f \Gamma_f^\ddagger)^2$, are functions of the quadratic lift and add no new invariant content. Additive shifts by multiples of the identity are excluded because they break the correspondence between the lift's kernel and the kernel of Γ_f, and hence corrupt the rank invariant. Conjugate contractions such as $\Gamma_f^* \Gamma_f^T$ transform under the conjugate frame representation and are anti-unitarily equivalent to the stated lifts, with identical spectra; they duplicate rather than extend the invariant content and are excluded as copies, not alternatives. All of this is relative to the declared structure per §4.5: an additional symmetry-fixed tensor, once declared by a later gate, may open further covariant combinations, whose admissibility then passes through HLF-13. Absent such structure, the quadratic lifts are minimal and unique in the stated sense, and the classification is exhaustive. ■

13. Theorem 2 — The Frame-Rotation Firewall Theorem

13.1 Statement

Theorem 2 — Frame-Rotation Firewall Theorem. Given an admissible closure map Γ_f, the following are not physical predictions unless additional invariant or symmetry-fixing structure is supplied per HLF-13:

1. raw entries of Γ_f;
2. arbitrary texture zeros of Γ_f;
3. diagonal, triangular, symmetric, or sparse coordinate forms of Γ_f;
4. eigenvectors of a single closure sector as absolute frame data;
5. phases removable by rephasing;
6. mixing angles inside degenerate eigenspaces;
7. CKM or PMNS matrices extracted from only one Hermitian lift;
8. more independent mixing parameters than the invariant count of Lemma 5′ permits.

Admissible prediction targets are instead:

1. spectra of Hermitian lifts;
2. singular values of closure maps;
3. invariant eigenvalue ratios and hierarchy exponents tied to them;
4. ranks and degeneracy patterns;
5. spectral projectors, as covariant classes;
6. relative projector overlaps between two lifts on a shared left frame;
7. commutator invariants;
8. rephasing-invariant quartet and CP quantities;
9. extension-declared neutrino relative-frame objects, with invariant class fixed by the extension type.

13.2 Proof

By HLF-4 and HLF-4′, the admissible frame rotations — independent on right bundles and unlocked left bundles, common on doublet-locked left bundles — change the coordinate representation of Γ_f while leaving the physics fixed. Lemma 2 shows that individual entries are not invariant. Lemma 3 shows that entry-level texture zeros are not invariant unless protected or invariantly re-expressed. Lemma 4 shows that a single sector carries invariant spectral content but no absolute mixing frame. Lemma 8 shows that degenerate eigenspace directions are arbitrary. Lemma 9 shows that raw phases are not CP predictions unless embedded in invariant combinations. Lemma 5′ shows that the invariant relative-frame content has a fixed parameter count, so any claim exceeding it contains convention.

By contrast, Theorem 1 shows that Hermitian-lift spectra are invariant. Projector overlaps, commutators, traces, determinants, ranks, degeneracy patterns, quartet invariants, and rephasing-invariant products are likewise invariant under allowed frame rotations, as established in §5.2, §8.5, §10.2, and §11.

Therefore the firewall separates admissible invariant content from inadmissible coordinate content. ■

---

14. Theorem 3 — The Relative-Frame Mixing Theorem

14.1 Statement

Theorem 3 — Relative-Frame Mixing Theorem. Mixing data in VERSF are relative-frame invariants, in the following precise sense.

1. A single closure sector determines no mixing observable: its full frame-invariant content is its singular spectrum.
2. Two left Hermitian lifts on a shared weak-generation frame — transported to the common generation factor $\mathcal{G}$ per the Shared Generation-Factor Lemma of §3.5 — determine a relative-frame matrix $V = U_1^\dagger U_2$, invariant under common frame rotation, whose physical content is its rephasing class after an ordering convention is fixed.
3. For n non-degenerate generations with Dirac-type residual freedom, that class carries exactly $n(n-1)/2$ angles and $(n-1)(n-2)/2$ phases; for three generations, three angles and one phase. Majorana-type extensions enlarge the phase count per §9.3.
4. The magnitude content is the projector geometry Tr(P_a^(1) P_b^(2)) = |V_ab|², and the CP-odd content is carried by the commutator invariants of the two lifts, per Proposition 1.
5. The relative-frame matrix is trivial, up to rephasing and permutation, exactly when the two lifts commute and their spectra are non-degenerate.
6. In the lepton sector, no relative-frame matrix exists in the minimal ledger; PMNS is admissible only after a neutrino-mass extension is declared, with invariant class fixed by the extension type.

14.2 Proof

Point 1 is Lemma 4. Point 2 follows from §8.3, HLF-4′, and Lemma 5 — the doublet lock is what forbids the independent left rotations that would trivialise V. Point 3 is Lemma 5′. Point 4 follows from §8.5 and Proposition 1, whose projector and commutator expressions are manifestly invariant. Point 5 is the corollary of §8.7 in both directions: commuting lifts with non-degenerate spectra are simultaneously diagonalisable, so V is diagonal up to permutation and rephasing; conversely a trivial rephasing class makes the lifts simultaneously diagonalisable, so they commute. Point 6 is Lemma 7. ■

---

15. Theorem 4 — Prediction-Admissibility for Later VERSF Gates

A later VERSF prediction is admissible only if it satisfies all five tests.

15.1 Test 1 — Ω-address target

The prediction must state which occupied slot or closure bundle it targets.

Examples:

$\mathfrak{U}_3$, $\mathfrak{E}_1$, $\mathscr{R}_u \to \mathscr{L}_u$, K_u^Q .

A particle label alone is insufficient.

15.2 Test 2 — Lift target

The prediction must say whether it targets:

- Γ_f as a closure map;
- K_f^L;
- K_f^R;
- singular values;
- eigenvalue ratios;
- spectral projectors;
- relative-frame data;
- commutator invariants;
- CP invariants.

15.3 Test 3 — Frame invariance

The predicted object must be invariant under the allowed frame rotations, or the paper must provide a symmetry or selection rule that fixes the frame, together with the residual frame group after fixing.

15.4 Test 4 — Degeneracy and ordering discipline

If eigenvalues are degenerate, the prediction must use projectors, not arbitrary eigenvectors. If eigenvalues are ordered, the ordering convention must be stated. Mixing claims must not exceed the invariant parameter count of Lemma 5′.

15.5 Test 5 — Extension discipline

If the prediction involves neutrino masses, PMNS, sterile states, vector-like fermions, Majorana sectors, or non-minimal Higgs fields, the extension ledger must be declared before the prediction is made, and the invariant class of any mixing claim must be stated for the declared extension type.

15.6 Statement

Theorem 4 — Prediction-Admissibility Theorem. A VERSF matrix prediction is admissible if and only if it passes Tests 1 through 5; equivalently, if and only if it targets an Ω-addressed Hermitian-lift invariant, a singular-spectrum invariant, or a relative-frame invariant between declared Hermitian lifts, expressed in the frame-invariant form required by HLF-13.

Proof. Necessity: a claim failing Test 1 violates HLF-1 and I7; failing Test 2 violates HLF-3 and HLF-5; failing Test 3 violates HLF-4 or HLF-4′ by Theorem 2; failing Test 4 violates HLF-9 or HLF-10 by Lemmas 8 and 5′; failing Test 5 violates HLF-12 by Lemma 7. Sufficiency: a claim passing all five tests targets an object shown invariant by Theorems 1–3, attached to a declared Ω-address, with degeneracy, ordering, rephasing, and extension freedom quotiented or declared. No frame convention can then alter the claim's truth value, which is the definition of admissibility used throughout this gate. ■

Status note. The sufficiency direction holds partly by construction: passing the five tests is what this gate defines admissibility to be. Theorem 4 is therefore graded in the closure certificate as closed by construction as the gate's admissibility rule, not as an independent structural result. The necessity direction, by contrast, is substantive, resting on Theorems 1–3 and the named premises.

---

16. Falsification Conditions

#	Failure	Consequence
F1	A mass prediction targets a raw matrix entry	Frame dependence
F2	A Yukawa eigenvalue is taken from a non-Hermitian raw matrix	Invalid spectral target
F3	A texture zero is claimed without symmetry or invariant protection	Coordinate artifact

#	Failure	Consequence
F4	A diagonal basis is treated as a physical derivation	Basis-choice error
F5	A single-sector diagonalising matrix is called CKM	Relative-frame failure
F6	PMNS is derived without a neutrino Hermitian lift	Extension violation
F7	Left and right Hermitian lifts are claimed to have different non-zero spectra	Singular-spectrum failure
F8	Degenerate eigenvectors are treated as physical	Degeneracy failure
F9	Rephasing-dependent phases are treated as CP predictions	Phase-convention failure
F10	Eigenvalue ordering is used without convention	Label ambiguity
F11	A prediction lacks an Ω-address target	Bridge failure
F12	Kinetic metric is non-positive or unspecified where needed	Hermitian lift invalid
F13	A matrix texture changes under allowed frame rotation	Texture not physical
F14	CKM is claimed when K_u^Q and K_d^Q are not both defined	Missing relative pair
F15	PMNS is claimed in the minimal massless-neutrino ledger	Minimal/extension confusion
F16	A commutator invariant is non-zero but simultaneous diagonalisation is claimed	Mixing inconsistency
F17	A commutator invariant is zero with non-degenerate spectra but non-trivial mixing is claimed	Relative-frame inconsistency
F18	An arbitrary flavour frame is treated as a VERSF-derived frame	Frame overclaim
F19	A raw phase is claimed as CP violation while removable by rephasing	CP overclaim
F20	A later hierarchy theorem predicts "generation 1" without tying ordering to Hermitian-lift eigenvalues	Generation-label drift
F21	A mixing claim delivers more independent content than the invariant class of its declared sector permits — three angles and one phase for Dirac-type three-generation classes; three angles and three phases for Majorana-type lepton extensions per §9.3	Parameter-count violation
F22	A CP claim is inconsistent with det(C_ud) computed from the declared lifts	Quartet-identity violation
F23	A structural constraint is asserted with no exhibited frame-invariant reformulation	Invariant-expressibility failure
F24	A Majorana-extension PMNS claim uses the Dirac rephasing class	Extension-class mismatch
F25	Independent left rotations are applied to weak-partner bundles of one doublet address	Doublet-locking violation
F26	A Majorana-extension gate imports SVD left/right machinery for a symmetric mass operator	Takagi-structure violation

#	Failure	Consequence
F27	A hierarchy exponent uses an untyped ε not declared as a frame-invariant Ω-addressed quantity	κ-ledger violation
F28	A commutator, projector overlap, or relative frame is formed between lifts without transport to the shared generation factor	Ill-typed comparison
F29	The negative spectral branch of the block operator $\mathcal{D}_\Gamma$ is treated as a physical mass, or $\mathcal{D}_\Gamma$ is treated as a primitive prediction target	Auxiliary-space overclaim

---

17. SF-1 Closure Certificate

SF-1 condition	Result
Occupied bundles defined	Closed
Closure map Γ defined	Closed
Frame rotations defined	Closed
Doublet-locked frame premise supplied	Closed
Shared generation-factor lemma supplied	Closed
Positive metric requirement stated	Closed
Metric adjoint supplied	Closed
Canonicalisation equivalence proved	Closed
Canonical Hermitian lift supplied	Closed
Left lift K_L supplied	Closed
Right lift K_R supplied	Closed
Lift positivity proved	Closed
Lift covariance proved	Closed
Singular-spectrum equivalence proved	Closed
Polar-factor identification proved	Closed
Minimality and uniqueness established, relative to declared structure	Closed
Block-Hermitian auxiliary firewall supplied	Closed
Raw matrix-entry firewall supplied	Closed
Texture-zero firewall supplied	Closed
Invariant-expressibility rule supplied	Closed
Diagonal-basis firewall supplied	Closed
Single-sector rotation firewall supplied	Closed
CKM typed as relative left-frame object	Closed
Relative-frame parameter count proved	Closed
Commutator–quartet identity derived in-corpus	Closed

SF-1 condition	Result
PMNS typed as extension-dependent relative left-frame object	Closed
Majorana/Dirac invariant-class distinction recorded	Closed
PMNS parameter counts per extension type recorded	Closed
Takagi structure of Majorana operators recorded	Closed
Rectangular-channel polar form supplied	Closed
Degeneracy firewall supplied	Closed
Rephasing/permutation firewall supplied	Closed
CP invariant target class supplied	Closed
Later-prediction admissibility rule supplied	Closed by construction as the gate's admissibility rule
Numerical mass values	Outside scope
Numerical Yukawa values	Outside scope
CKM numerical entries	Outside scope
PMNS numerical entries	Outside scope
Neutrino mass mechanism	Extension gate
Higgs vev	Outside scope

SF-1 closure grade: closed as a Hermitian-lift and frame-rotation firewall theorem. It fixes the admissible matrix objects that later VERSF Standard Model derivations must target, the exact invariant parameter budget available to mixing claims, and the invariant class available to CP claims.

---

18. Conclusion

SC-2 fixed where Standard Model matter sits inside VERSF. SF-1 fixes what may act on those occupied slots.

The result is simple but important: a raw matrix is not yet physics. A matrix becomes physically admissible only through its Hermitian lift, its singular spectrum, or its relative-frame relation to another Hermitian lift on a shared weak-generation frame.

For a closure map Γ, the minimal physical lifts are

$K_L = \Gamma \Gamma^\dagger$, $K_R = \Gamma^\dagger \Gamma$.

These lifts are Hermitian, positive, frame-covariant, and spectrum-invariant. They are the squared polar factors of the closure map, so they carry exactly the frame-invariant magnitude content and nothing else. Their eigenvalues are the squared singular values of the closure map. These are valid targets for later mass and hierarchy derivations.

CKM is not an arbitrary rotation. It is the relative orientation between the up-sector and down-sector left Hermitian lifts on the shared quark weak-generation frame, physical only as a rephasing class, and budgeted at exactly three angles and one phase for three non-degenerate generations. Its CP-odd content is fixed by the commutator of the two lifts through the quartet identity. The comparison exists at all because HLF-4′ locks the two left frames to the shared doublet address — the lock is a named premise with its own falsifier, and it is the structural reason mixing is physical.

PMNS is not available in the minimal massless-neutrino ledger. It becomes meaningful only after a neutrino-mass extension supplies a neutrino Hermitian lift on the shared lepton weak-generation frame, and its invariant class depends on whether the extension is Dirac-type or Majorana-type.

Texture zeros, diagonal forms, raw entries, and isolated matrix phases are firewalled away unless re-expressed as frame-invariant conditions or protected by symmetry, selection rule, rank condition, or projector relation.

In one sentence:

The minimal Hermitian lift turns VERSF mass and mixing claims from coordinate-dependent matrix statements into frame-invariant spectral and relative-frame predictions, with a fixed invariant parameter budget, ensuring that later Standard Model derivations cannot mistake a basis choice for physics.

Appendix A — Minimal Algebraic Skeleton

Closure map:

$\Gamma : \mathscr{R} \to \mathscr{L}$.

Canonical frame rotations:

Γ ↦ U_L† Γ U_R .

Hermitian lifts:

K_L = Γ Γ† , K_R = Γ† Γ .

Transformation:

$$K_L \mapsto U_L^\dagger K_L U_L , \quad K_R \mapsto U_R^\dagger K_R U_R .$$

Polar decomposition:

$\Gamma = P W$, $P = \sqrt{\Gamma \Gamma^\dagger}$ unique , W unitary (square channels; canonical partial isometry in general) , $K_L = P^2$.

Singular-value form:

$$\Gamma = U_L \Sigma U_R^\dagger ,$$

$$K_L = U_L \Sigma^2 U_L^\dagger , \quad K_R = U_R \Sigma^2 U_R^\dagger .$$

Physical spectral targets:

$$\lambda_i(K_L) = \lambda_i(K_R) = \sigma_i^2 .$$

Yukawa-to-mass relation after electroweak symmetry breaking:

$$m_i = (v / \sqrt{2}) \, \sigma_i(Y_f) .$$

Shared generation factor:

$Q_L \cong \mathcal{G}_Q \otimes \mathbb{C}^2_{weak}$, $\mathscr{L}_u = \mathcal{G}_Q \otimes e_+$, $\mathscr{L}_d = \mathcal{G}_Q \otimes e_-$, both left lifts transported to $\mathrm{End}(\mathcal{G}_Q)$.

Canonicalisation:

$$\hat{\Gamma} = G_L^{(1/2)} \Gamma G_R^{(-1/2)} , \quad \hat{\Gamma} \hat{\Gamma}^\dagger = G_L^{(1/2)} (\Gamma \Gamma^\ddagger) G_L^{(-1/2)} .$$

Block-Hermitian auxiliary, not a primitive target:

$$\mathcal{D}_\Gamma = \begin{pmatrix} 0 & \Gamma \\ \Gamma^\ddagger & 0 \end{pmatrix} , \quad \mathcal{D}_\Gamma^2 = \begin{pmatrix} \Gamma \Gamma^\ddagger & 0 \\ 0 & \Gamma^\ddagger \Gamma \end{pmatrix} .$$

Quark relative frame:

$V_{CKM} = U_{uL}^\dagger U_{dL}$, physical as the class $\Phi_u^\dagger V \Phi_d$.

Parameter count, n non-degenerate generations:

$$n(n-1)/2 \text{ angles} , \quad (n-1)(n-2)/2 \text{ phases} .$$

Projector magnitude:

$$|V_{ab}|^2 = \mathrm{Tr}(P_a^u P_b^d) .$$

Commutator:

$C_{ud} = [K_u^Q, K_d^Q]$.

Quartet identity, three generations:

$\det(C_{ud}) = 2i\, J \cdot \Pi_{(a<a')} (\lambda_{a'}^u - \lambda_a^u) \cdot \Pi_{(b<b')} (\lambda_{b'}^d - \lambda_b^d)$,

$J = \mathrm{Im}(\, V_{11}\, V_{22}\, V_{12}\, V_{21}\,)$.

If $C_{ud} = 0$ and spectra are non-degenerate, the lifts are simultaneously diagonalisable and mixing is trivial.

A.1 Derivation of the quartet identity

Work in the frame that diagonalises K_u^Q, admissible as a common left rotation per HLF-4′:

$K_u^Q = D_u = \mathrm{diag}(\lambda_1^u, \lambda_2^u, \lambda_3^u)$, $K_d^Q = M = V D_d V^\dagger$, $M_{ij} = \Sigma_b\, \lambda_b^d\, V_{ib}\, V^*_{jb}$.

Step 1 — Determinant of the commutator. The commutator has entries

$C_{ij} = (\lambda_i^u - \lambda_j^u)\, M_{ij}$,

so its diagonal vanishes. For a three-by-three matrix with zero diagonal, only the two cyclic permutations contribute to the determinant:

$\det C = C_{12}\, C_{23}\, C_{31} + C_{13}\, C_{21}\, C_{32}$.

Step 2 — Extracting the up-sector Vandermonde. The eigenvalue prefactors of the two cyclic products are Δ^u and $-\Delta^u$ respectively, where

$\Delta^u = (\lambda_1^u - \lambda_2^u)(\lambda_2^u - \lambda_3^u)(\lambda_3^u - \lambda_1^u) = \Pi_{(a<a')} (\lambda_{a'}^u - \lambda_a^u)$.

Since M is Hermitian, $M_{13}\, M_{21}\, M_{32} = (M_{12}\, M_{23}\, M_{31})^*$, hence

$\det C = \Delta^u \cdot 2i\, \mathrm{Im}(M_{12}\, M_{23}\, M_{31})$.

Step 3 — Vandermonde divisibility in the down sector. $\mathrm{Im}(M_{12}\, M_{23}\, M_{31})$ is a real cubic polynomial in the λ^d. If any two coincide, say $\lambda_1^d = \lambda_2^d = \mu$, then by unitarity $M = \mu\, \mathbb{1} + (\lambda_3^d - \mu)\, w\, w^\dagger$ with w the third column of V, and

$M_{12}\, M_{23}\, M_{31} = (\lambda_3^d - \mu)^3\, |w_1|^2\, |w_2|^2\, |w_3|^2$,

which is real, so the imaginary part vanishes — and symmetrically for the other two coincidence pairs. The polynomial is therefore divisible by the cubic Vandermonde Δ^d, and by degree matching

$\mathrm{Im}(M_{12} M_{23} M_{31}) = c(V) \cdot \Delta^d$,

with c(V) independent of the λ^d.

Step 4 — Unitarity fixes the quartet structure. For a unitary V, the quartets Q_(ij,ab) = V_ia V_jb V_*ib V*_ja satisfy, from Σ_b V_jb V*_ib = 0 for i ≠ j, the row-pair sum rule Σ_b Im Q = 0, and symmetrically the column-pair sum rule Σ_j Im Q = 0. For a three-by-three unitary these constraints force

Im(V_ia V_jb V_*ib Vja) = J · Σ*(k,c) ε_ijk ε_abc ,

for a single real $J = \mathrm{Im}(V_{11} V_{22} V^*_{12} V^*_{21})$.

Step 5 — Coefficient evaluation. Set $\lambda_3^d = 0$ and compare coefficients of $\lambda_1^{d2} \lambda_2^d$. In $M_{12} M_{23} M_{31}$ this coefficient collects three terms; by Step 4 their imaginary parts are $-J|V_{31}|^2$, $-J|V_{11}|^2$, and $-J|V_{21}|^2$, summing by column unitarity to −J. In Δ^d at $\lambda_3^d = 0$, which equals $-\lambda_1^{d2} \lambda_2^d + \lambda_1^d \lambda_2^{d2}$, the same coefficient is −1. Hence c(V) = J, and

$\det C\_ud = 2i\, J\, \Delta^u \Delta^d$. ■

Appendix B — Channel Lift Table

Sector	Closure map	Higgs route	Left lift	Right lift	Later physical target
Up quarks	Γ_u : ℛ_u → ℒ_u	H̃	K_u^Q = Γ_u Γ_u†	K_u^U = Γ_u† Γ_u	u, c, t singular spectrum
Down quarks	Γ_d : ℛ_d → ℒ_d	H	K_d^Q = Γ_d Γ_d†	K_d^D = Γ_d† Γ_d	d, s, b singular spectrum
Charged leptons	Γ_e : ℛ_e → ℒ_e	H	K_e^L = Γ_e Γ_e†	K_e^E = Γ_e† Γ_e	e, μ, τ singular spectrum
Neutrinos, minimal	none	none	none	none	no minimal mass lift
Neutrinos, extension	Γ_ν or effective operator	extension-dependent	K_ν^L	Dirac: Γ_ν† Γ_ν; Majorana: structurally absent (Takagi factorisation, $m\_\nu = U D U^T$, single-sided frame freedom)	neutrino spectrum and PMNS, invariant class per extension type

---

Appendix C — Referee Objections and Replies

Objection 1 — Is this just the singular-value decomposition?

Reply. Mathematically, the Hermitian lift uses the same spectral discipline underlying singular-value decomposition and polar factorisation. The contribution of SF-1 is not to invent SVD; it is to install SVD/Hermitian-lift discipline as a compulsory VERSF bridge rule, with named premises, falsifiers, an invariant parameter budget, and an admissibility theorem. Later VERSF mass and mixing claims must target invariant spectral or relative-frame data, not arbitrary matrix entries.

Objection 2 — Why not use eigenvalues of the Yukawa matrix directly?

Reply. A generic Yukawa matrix is complex and non-Hermitian. Its eigenvalues are not the physical fermion masses, are not invariant under independent left/right frame rotations, and for a non-normal matrix do not coincide with its singular values. The physical magnitudes are singular values, equivalently eigenvalues of the Hermitian lifts.

Objection 3 — Does this forbid texture models?

Reply. No. It forbids unprotected texture claims. A texture is admissible if a symmetry, Ω-address selection rule, rank condition, invariant relation, or frame-fixing principle makes it physical — and HLF-13 gives the constructive test: exhibit the frame-invariant functional whose vanishing the texture expresses. A naked matrix-entry zero is not enough.

Objection 4 — Does this derive CKM?

Reply. No. SF-1 defines what CKM is allowed to be in VERSF: a relative left-frame rephasing class between the up-sector and down-sector Hermitian lifts on the shared quark weak-generation frame, carrying exactly three angles and one phase. Numerical CKM entries require a later theorem, which now has a fixed budget of four invariants to deliver.

Objection 5 — Does this derive PMNS?

Reply. No. PMNS requires a neutrino Hermitian lift. The minimal SC-2 ledger has left-handed neutrino components but no minimal right-handed sterile closure slot or declared Majorana mechanism. PMNS is therefore extension-dependent, and even its invariant class — the number of physical phases — depends on the declared extension type.

Objection 6 — Why are there left and right lifts if they have the same non-zero spectrum?

Reply. They act on different occupied bundles. The left lift determines the left eigenframe relevant for charged-current mixing. The right lift determines the right-terminal spectral frame. Their non-zero spectra agree, but their frame roles are different.

Objection 7 — Are phases physical?

Reply. Only rephasing-invariant phases are physical. A phase in a raw matrix entry can usually be moved by field redefinitions. CP claims must be expressed through invariant products, commutators, projectors, or equivalent rephasing-safe objects. In the three-generation quark sector the entire CP-odd content is fixed by the commutator–quartet identity of Proposition 1.

Objection 8 — What happens if two masses are equal?

Reply. Then individual eigenvectors inside the degenerate subspace are not physical. Only the projector onto the degenerate subspace is invariant. Mixing angles inside that subspace require additional structure, and the degeneracy-safe replacement for $|V_{ab}|^2$ is the subspace overlap $\mathrm{Tr}(P_S P_T)$. Note also that exact degeneracy in either sector forces the CP-odd determinant of Proposition 1 to vanish.

Objection 9 — Does the Hermitian lift depend on choosing canonical kinetic terms?

Reply. The canonical formula uses †. If the kinetic metric is non-trivial but positive, the correct object uses the metric adjoint ‡. One may either canonicalise the kinetic terms first or use the metric-adjoint form directly; the resulting spectra agree because the two procedures are related by a positive similarity transformation.

Objection 10 — Is the minimality claim of Theorem 1 merely aesthetic?

Reply. No. It is a classification result. Linear constructions fail to be endomorphisms of a single bundle. Quadratic covariant contractions are exhausted by $\Gamma \Gamma^‡$ and $\Gamma^‡ \Gamma$ up to positive scale. Higher powers are functions of the quadratic lifts and add no invariant content. Identity shifts corrupt the rank invariant. The quadratic lift is therefore the unique minimal admissible object, not a stylistic preference.

Objection 11 — Does the parameter count assume anything hidden?

Reply. It assumes non-degenerate spectra and a stated ordering convention, both flagged as premises. Under degeneracy the residual frame group is larger than the diagonal torus, the count changes, and the projector formulation of §10.2 takes over. The count also assumes a Dirac-type rephasing class; a Majorana extension reduces the neutrino-side torus to signs and enlarges the invariant phase content, as recorded in §9.3.

Objection 12 — Doesn't HLF-4 let me diagonalise both quark lifts at once and dissolve CKM?

Reply. Not under the admissible frame group. HLF-4′ types $\mathscr{L}_u$ and $\mathscr{L}_d$ as weak components of the single occupied doublet address Q_L, carrying one shared generation frame and admitting only a common left rotation. Independent left rotations of weak partners are not frame changes at all, because admissible frames must respect the doublet address structure. CKM survives precisely because the doublet lock removes the rotations that would trivialise it — this is the deepest structural fact in the gate, and it is carried by a named premise with falsifier F25, not by an unstated assumption.

Objection 13 — Isn't the block operator $\mathcal{D}_\Gamma$ a linear Hermitian alternative to the quadratic lift?

Reply. It is Hermitian but not positive, and it acts on the doubled auxiliary space $\mathscr{L} \oplus \mathscr{R}$, which is not an occupied Ω-address bundle — no occupancy assignment supplies that space. Its square is block diagonal in the two quadratic lifts, so its invariant positive content is exactly theirs, and its $\pm\sigma_i$ spectrum adds only a doubling artifact in the negative branch. It is admissible as an auxiliary representation and inadmissible as a primitive prediction target; the firewall is §4.5 and the falsifier is F29.

Objection 14 — What is the one-sentence result?

Reply.

VERSF matrix predictions become admissible only after raw closure maps are lifted to Hermitian, positive, frame-covariant operators, so masses are spectral invariants, mixings are relative-frame rephasing classes with a fixed parameter budget, CP resides in commutator–quartet invariants, and arbitrary basis choices are not physics.

Final One-Sentence Theorem

Given the SC-2 world-to-framework occupancy map, any admissible VERSF mass, Yukawa, hierarchy, CKM, PMNS, or CP prediction must pass through the minimal Hermitian lift of an Ω-addressed closure map — the squared polar factor of that map —

targeting only singular spectra, Hermitian-lift eigenvalues, projector geometry, commutator and quartet invariants, or rephasing-safe relative-frame classes within the invariant parameter budget; therefore raw matrix entries, arbitrary texture patterns, diagonal-basis choices, and unconstrained frame rotations are firewalled from physical prediction.