The Final Eight Papers of the VERSF Standard Model Programme

From Quantum Gauge Architecture to a Complete, Renormalised and Auditable Derivation

Keith Taylor VERSF Theoretical Physics Programme

General Reader Summary

The VERSF Standard Model programme has reached an unusual point in the life of a physical theory.

Most of the major structures that conventional particle physics treats as inputs are now situated inside a VERSF theorem architecture. The programme contains gate theorems on chirality, gauge structure, the matter census, the Higgs-radial carrier, Yukawa operators, quark and lepton mass architecture, CKM and PMNS frame geometry, QCD running and confinement, the strong phase, and quantum gauge consistency.

The most recently completed gate, **The Quantum Gauge Completion of the VERSF Standard Model**, closes a milestone in that sequence. It establishes how the classically derived VERSF Standard Model is embedded in the local perturbative machinery of a quantum gauge theory: gauge-orbit quotienting, Faddeev–Popov determinants, ghosts, BRST and BV symmetry, anomaly cancellation, Slavnov–Taylor identities, renormalisation, running couplings, physical-state cohomology and perturbative unitarity. No further paper is needed whose task is merely to explain how the VERSF Standard Model is quantised.

That same paper is equally explicit about the boundary of what it achieves. It does not construct the microscopic substrate measure, solve the global Gribov problem, build the full quotient-compatible chiral determinant measure, establish the complete faithful-group global anomaly classification, or evaluate the numerical boundary values of the Standard Model parameters. Its perturbative quantum architecture is conditionally closed; its global nonperturbative structure and its numerical completion remain open.

That boundary defines the final phase of the programme.

The remaining work is no longer principally about discovering the correct **form** of the answer. In almost every sector, the form is already fixed by an existing gate theorem. The remaining task is to **calculate** the actual VERSF objects, assemble them at one declared matching scale, transport them through renormalisation and thresholds, and audit the result without inserting measured Standard Model values anywhere in the chain.

This paper sets out the eight programme-level papers required to do that honestly, states what each must deliver, gives each an explicit pass/fail closure test, and fixes the dependency order in which the work should proceed.

The Non-Circularity Discipline

Every closure test below invokes one shared discipline, stated here once:

> **No measured Standard Model value may enter any calculation upstream of the point at which that value is claimed as an output.** Observed couplings, masses, mixing angles, phases and scales may be consulted only *after* the corresponding frozen VERSF object has been calculated, published at its matching scale, and transported. A quantity selected, tuned or normalised against experiment is an import, not a derivation, wherever it subsequently appears.

A paper that satisfies its own closure test under this discipline is said to be **numerically closed**. A paper that fixes the form of an object but not its value is **structurally closed**. The purpose of the final phase is to convert every structurally closed sector into a numerically closed one — or to record precisely where the conversion fails.

The Eight Papers at a Glance

#	Paper	Sector	Consumes	Produces
1	Electroweak Vacuum and Higgs-Potential Completion	Scalar	Higgs-radial gate (HR-1)	v_{cl}, m_H, λ, μ_H^2
2	Neutral-Fermion Branch and Carrier-Census Completion	Neutrino census	Neutral-completion rulebook (NMCC-1), PMNS frame geometry (PMNS-1)	Frozen neutral-carrier ledger; typed completion operator
3	Microscopic Closure-Hessian and Full Yukawa Evaluation	Flavour	Yukawa construction gate (CMSY-1); Paper 2	Y_u, Y_d, Y_e, Y_ν (or neutral blocks)
4	Gauge-Response and Coupling Evaluation	Gauge couplings	Gauge action and coupling gate (GAC-1)	$Z_3(\mu\star)$, $Z_2(\mu\star)$, $Z_q(\mu\star)$; g_s, g, g'
5	Non-Perturbative QCD Scale and Triality-Gap Completion	Strong sector	Consolidated QCD gate (QCR-1); Paper 4	χ_C, Λ_{QCD}, quantum triality gap
6	Substrate Cohomology, Global Gauge Measure, and Strong-Phase Trivialisation	Global measure	Quantum Gauge Completion	μ_{cont}, global anomaly audit, $\bar{\theta}$
7	Renormalised Standard Model Parameter-Closure	Assembly	Papers 1–6	One transported observable set

#	Paper	Sector	Consumes	Produces
8	Final Standard Model Derivation Audit	Audit	Papers 1–7	Graded row-by-row verdict

1. The Electroweak Vacuum and Higgs-Potential Completion Theorem in VERSF

Why this paper is still required

The existing Higgs-radial gate answers the prior structural question:

> Why must VERSF contain a Higgs-like scalar radial mode at all?

It identifies the Higgs-radial degree of freedom with the normal return-stiffness mode of the closure interface. That derivation matters: it prevents the scalar field from being inserted merely because the Standard Model needs one.

The same gate explicitly separates this structural result from the numerical Higgs-sector problem. It does not derive the microscopic closure-interface potential, its kinetic normalisation, the electroweak vacuum scale, the physical Higgs mass or the quartic self-coupling. Those objects are the load of the present paper.

What the paper must deliver

The paper must construct the scalar effective action from VERSF substrate dynamics:

$\Gamma\_scalar = \int d^4x \, [\, \frac{1}{2} Z\_H(\varphi)\, (\partial\varphi)^2 - V\_CI(\varphi) + \cdots]$

where V_CI is the closure-interface potential and Z_H the substrate-derived kinetic normalisation. From this action it must derive, rather than import,

$v\_cl, m\_H, \lambda, \mu\_H^2$,

together with the vacuum stability conditions and the broken-phase relations that connect the scalar vacuum to the gauge and fermion sectors.

The quantum calculation must be **Nielsen-consistent**: gauge-dependent off-shell features of an effective potential cannot be presented as physical predictions. The physical outputs must be defined through gauge-consistent extrema, curvatures, pole positions and matching conditions — quantities on which the Nielsen identities guarantee gauge-parameter independence.

Closure test

Pass: the frozen VERSF substrate returns the electroweak vacuum and Higgs-sector parameters under the non-circularity discipline — in particular, without $v \simeq 246$ GeV or $m\_H \simeq 125$ GeV appearing anywhere upstream of the output.

Fail: any Higgs-sector value is fixed, bounded or normalised by reference to its observed counterpart before the calculation closes.

---

2. The Neutral-Fermion Branch and Carrier-Census Completion Theorem in VERSF

Why this paper is essential

The neutrino sector cannot be completed by writing the most general mass matrix and leaving every branch open.

The existing neutral-fermion gate supplies the complete mathematical rulebook for Dirac, Majorana, pseudo-Dirac and seesaw completion. It correctly distinguishes active carriers, sterile carriers, charge-conjugate carriers, Takagi masses, lepton-number symmetry and the exact Schur-complement structure of the seesaw. But it deliberately leaves open the existence and multiplicity of right-handed neutral carriers, the physical completion branch, the blocks M_L, M_D, M_R, the absolute spectrum and the ordering.

The PMNS gate likewise proves that lepton mixing is the relative spectral geometry of the charged-lepton and neutrino frames, but leaves the active–sterile census, the physical mass ordering and the electron–muon versus electron–tau attachment branch unresolved.

A rulebook with every branch open is not a prediction. This paper must select the branch — from the substrate, not from the data.

What the paper must deliver

It must determine from the frozen substrate:

- whether right-handed neutral carriers exist;
- how many such carriers exist;
- whether exact or broken B−L structure is present;
- which of M_L, M_D and M_R are nonzero;
- whether the physical spectrum is Dirac, Majorana, pseudo-Dirac, seesaw or mixed;
- whether the three light states exhaust the active carrier;
- the physical neutrino mass ordering;
- the physical atmospheric-octant and CP-sign branch.

The output must be a **frozen neutral-carrier ledger** and a **uniquely typed neutral completion operator** — the object every downstream paper consumes wherever a neutrino Yukawa matrix would otherwise be assumed.

Closure test

Pass: the branch is selected before oscillation data, cosmological mass bounds or neutrinoless-double-beta-decay results are consulted, and the selection is traceable to named substrate premises.

Fail: the branch is chosen because it fits the known neutrino data. That is model selection, not a VERSF derivation, however well it fits.

---

3. The Microscopic Closure-Hessian and Full Yukawa Evaluation Theorem in VERSF

Why this is the central remaining calculation

The Closure-Operator Yukawa Construction Theorem has already fixed the leading operator form:

$Y_f^c = \mathcal{E}_{fL}^\dagger \hat{C}_Y \mathcal{E}_{fR}$, with $\hat{C}_Y = P_Y \hat{H}_{cl} P_Y$,

where $\hat{H}_{cl}$ is the closure Hessian, P_Y the scalar-readout projector, and $\mathcal{E}_{fL}$, $\mathcal{E}_{fR}$ the chiral carrier embeddings.

That gate also closes the obvious loopholes in advance: the response function cannot be selected freely, the embeddings cannot be fitted, phases must be audited through rephasing invariants, and texture zeros must arise from projector-level support rather than coordinate choice.

But the gate is explicit that the numerical closure Hessian, the scalar-readout projector and the chiral embeddings remain load-bearing open calculations. It derives the **form** of the Yukawa answer, not the four matrices. Every flavour observable in the programme — every mass ratio, every mixing angle, every CP phase — waits on this evaluation. That is why this paper carries more of the programme's remaining weight than any other.

What the paper must deliver

It must calculate, from one frozen VERSF substrate:

$\hat{H}_{cl}$, P_Y, $\mathcal{E}_{uL/R}$, $\mathcal{E}_{dL/R}$, $\mathcal{E}_{eL/R}$, $\mathcal{E}_{\nu L/R}$,

with the neutral embeddings determined by the carrier census frozen in Paper 2 — not assumed independently.

It must then return

Y_u, Y_d, Y_e, Y_ν

or, where the physical neutrino branch is not purely Dirac, the corresponding neutral symmetric completion blocks.

The calculation must produce the singular spectra, ranks, physical left and right projectors, determinant phases, CKM relative frame, charged-lepton frame and PMNS-compatible neutrino frame — as computed outputs of the frozen operators, with no residual gauge of convention hiding a fitted quantity.

Closure test

Pass: all four matrices (or their neutral completion blocks) are calculated and published before the observed masses, mixing angles and CP phases are consulted.

Fail: any entry, texture, phase or hierarchy is adjusted against flavour data at any stage.

This is the paper where the programme must move from

“the answer must have this form”

to

“the frozen substrate returns these operators.”

4. The Gauge-Response and Coupling Evaluation Theorem in VERSF

Why the gauge-coupling gate is not yet the coupling calculation

The Gauge Action and Coupling Theorem derives the correct gauge-response architecture. It identifies the gauge kinetic operator as the auxiliary-reduced curvature Hessian,

$\mathcal{K}_g = H_{FF} - H_{F\eta} H_{\eta\eta}^{-1} H_{\eta F}$,

and establishes

$g_s = Z_3^{-1/2}$, $g = Z_2^{-1/2}$, $g' = 6 Z_q^{-1/2}$

in the declared primitive hypercharge convention. It also proves coupling universality: the same canonical field normalisation controls matter vertices and non-Abelian self-interactions, so a single response coefficient per factor fixes every vertex in that factor.

What remains open is the construction and evaluation of the actual factorwise response coefficients Z_3, Z_2 and Z_q, together with the matching scale, the primitive current normalisation and the scheme bridge.

The Quantum Gauge Completion paper supplies the perturbative transport once those boundary values exist — it recovers the Standard Model one-loop coefficients and the all-order Slavnov–Taylor architecture — but it does not, and cannot, supply the boundary values themselves. Transport without boundary data moves nothing.

What the paper must deliver

It must construct the holonomy-response

functional on the frozen substrate and calculate

$Z_3(\mu\star)$, $Z_2(\mu\star)$, $Z\_q(\mu\star)$.

It must also derive:

- the matching scale $\mu\star$ itself, from substrate structure rather than convenience; - the primitive current normalisation; - the auxiliary-mode census entering each reduced Hessian; - the finite scheme bridge to a standard perturbative scheme; - the threshold ledger; - the conversion chain to g_s, g, g', e, α and $\sin^2\theta\_W$.

Closure test

Pass: the three coupling values are published at the frozen matching scale $\mu\star$ before being run to any experimentally familiar scale, and no measured coupling normalises the response Hessian or selects μ⋆.

Fail: the matching scale or any response coefficient is reverse-engineered from the known low-energy couplings.

5. The Non-Perturbative QCD Scale and Triality-Gap Completion Theorem in VERSF

Why QCD still needs a final nonperturbative paper

The consolidated QCD gate has made major progress. It derives the colour-holonomy map, reduces the colour response to the scalar form

$Z_3 = \chi\_C/8 - b^\dagger C^{-1} b$,

and obtains a symbolic three-flavour strong scale after threshold matching. It also provides a conditional triality-gap route to confinement and — correctly — keeps asymptotic freedom and confinement as logically separate claims.

But that gate explicitly leaves open: the numerical colour susceptibility χ_C, the auxiliary adjoint census entering b and C, the global centre-sensitive lift, the primitive flux normalisation, the bath areal coefficient B_τ, the quantum promotion of the classical slice gap, the longitudinal slicing inequality, and the uniform continuum gap. Each of these is load-bearing for the claim that VERSF *derives* — rather than gestures at — confinement and the strong scale.

What the paper must deliver

Running side. In a declared scheme, using quark-threshold values supplied by the completed Yukawa and parameter-closure programme, it must calculate

$\chi\_C(\mu\star)$, b, C, $Z_3(\mu\star)$, Λ_QCD.

Confinement side. It must derive:

- the global centre-sensitive bath lift;
- the primitive triality-flux normalisation;
- the complete quantum transverse bath functional;
- coercivity, compactness and vacuum uniqueness for that functional;
- the longitudinal slicing inequality;
- the quantum nonzero-triality gap;
- the continuum and infinite-volume uniform limit of that gap;
- the physical string tension, or its full-QCD string-breaking replacement.

Closure test

Pass: a quantum, regulator-controlled, nonzero-triality-sector gap that survives the continuum and infinite-volume limits, plus a numerically evaluated Λ_QCD under the non-circularity discipline.

Fail: asymptotic freedom offered as a proxy for confinement. A negative beta function is not sufficient and is not the claim.

---

6. The Substrate Cohomology, Global Gauge Measure, and Strong-Phase Trivialisation Theorem in VERSF

Why the Quantum Gauge Completion paper does not eliminate this paper

The Quantum Gauge Completion paper closes the local perturbative architecture conditionally. It derives the Faddeev–Popov, ghost, BRST, BV, anomaly, renormalisation and physical-cohomology rulebook.

It is equally explicit that the following remain open: the microscopic measure pushforward, the global quotient-compatible bundle sum, the determinant-line topology, Gribov control, the nonperturbative chiral measure, the global anomaly classification on the faithful group, reflection positivity, and the nonperturbative physical Hilbert space. Its own conclusion identifies this global measure construction as the next programme-level target.

This paper should also close the unresolved physical strong-phase line. Strong CP is not an unrelated numerical curiosity: it lives on exactly the combined determinant line whose global structure this paper must construct, and it should be settled there.

What the paper must deliver

It must construct the global quantum measure from the substrate,

$\mu\_cont = \mathfrak{Q}_* \mu\_sub$,

as the pushforward of the substrate measure onto the complete faithful-group configuration space.

It must derive:

- gauge-basicity of the pushed-forward measure;
- the complete faithful-quotient bundle census;
- the determinant-line and bordism anomaly audit;
- the nonperturbative chiral measure;
- a global treatment of Gribov copies, or a gauge-invariant substitute construction;
- reflection positivity and the physical Hilbert-space reconstruction;
- the combined topological and quark determinant line.

The strong-phase result must be stated on the invariant combined line,

$\bar{\theta} = \theta_3 + \arg \det(M\_u\ M\_d)$,

and the paper must calculate whether the frozen substrate yields a positive real trivialisation of that line, a nonzero invariant phase, or no consistent trivialisation at all. Each of the three outcomes is a result; only silence is not.

Closure test

Pass: the result is global and quotient-compatible on the faithful group, and $\bar{\theta}$ is evaluated (or its obstruction exhibited) on the invariant line.

Fail: local anomaly cancellation or perturbative BRST closure reused as a substitute for the global measure calculation.

---

7. The Renormalised Standard Model Parameter-Closure Theorem in VERSF

The role of the final assembly paper

The first six papers calculate the missing sector objects. This paper assembles them into one Standard Model at one declared matching scale.

Its principal risk is well defined in advance: it must not become the hiding place in which unfinished microscopic calculations are quietly replaced with measured boundary values. Assembly is transport, not rescue.

The frozen parameter set takes the form

$\mathcal{P}\_SM(\mu\star) = \{ g\_s, g, g', \mu\_H^2, \lambda, Y\_u, Y\_d, Y\_e, \mathcal{B}\_\nu, \bar{\theta} \}$, where $\mathcal{B}\_\nu$ denotes the complete neutral completion object appropriate to the branch selected in Paper 2 — not an assumed Dirac Yukawa matrix where the substrate has selected a Majorana or mixed completion.

What the paper must deliver

Starting from the one frozen boundary set, it must perform the full transport chain:

$\mathcal{P}$_SM(μ⋆) → RG evolution → threshold matching → electroweak breaking → running masses → poles and physical observables.

The outputs must include:

- m_W, m_Z and the Higgs pole;
- all charged-fermion running and pole masses;
- the CKM matrix and its invariant CP sector;
- the PMNS matrix, the neutrino masses and any predicted nonunitarity;
- the gauge couplings at standard comparison scales;
- the QCD scale and strong-sector matching outputs;
- the physical strong-phase invariant;
- explicit uncertainty and truncation budgets for every transported quantity;
- the cross-sector consistency identities that one shared substrate implies.

Closure test

Pass: every sector uses the same frozen substrate, matching scale, conventions and threshold ledger, and every output carries a stated uncertainty budget.

Fail: a value imported in one sector and presented as a prediction in another — the single most common way a derivation programme silently becomes a fit.

---

8. The Final Standard Model Derivation Audit

What the audit is for

The final paper derives no new Standard Model structure. Its sole job is to determine whether the complete chain has actually succeeded — and to say so in a form a hostile referee can check row by row.

It must trace every physical object through the full chain:

VERSF primitive → carrier → operator → normalisation → renormalised parameter → observable.

What the audit must contain

Every Standard Model row is graded as exactly one of:

- **fully derived and numerically evaluated**;
- **derived conditionally** on a named programme premise;
- **reduced to a calculation** but not evaluated;
- **inherited** from external physics;
- **contradicted or falsified**;
- **outside the declared scope**.

The audit must test:

- completeness of the particle and interaction census;
- absence of hidden Standard Model insertion anywhere in the chain;
- consistency of conventions across all papers;
- reuse of the same substrate objects across independent sectors;
- anomaly and global-measure closure at the claimed level;
- renormalisation and threshold consistency;
- numerical agreement and disagreement with experiment, stated symmetrically;
- the count and identity of any remaining free parameters;
- whether any extra low-energy states or interactions are predicted.

Closure test

Pass: no load-bearing Standard Model row remains disguised as "to be calculated later," and every row carries exactly one grade.

A failed or partial audit remains scientifically valuable, provided the failed gate is identified precisely and the programme claims no more than the graded ledger supports. The audit's obligation is accuracy, not vindication.

Recommended Order of Work

The remaining papers need not be written serially. The efficient order is:

Phase I — Independent foundational completion. The Higgs-potential paper (1), the neutral-carrier paper (2), the gauge-response paper (4) and the global-measure paper (6) can proceed in parallel: none consumes another's frozen output.

Phase II — Microscopic evaluation. The full Yukawa evaluation (3) begins once the neutral carrier census (2) is frozen. The QCD numerical and confinement completion (5) proceeds once the colour response (4) and the global centre lift (6) are available.

Phase III — Common-scale assembly. The Parameter-Closure paper (7) consumes the outputs of all six preceding papers and produces one transported observable set.

Phase IV — Audit. The Final Derivation Audit (8) is written last and contains no new derivational rescue work.

The dependency structure is therefore:

Papers 1, 2, 4, 6 → Papers 3, 5 → Paper 7 → Paper 8

with limited overlap where the calculations permit it.

Why These Papers Cannot Be Replaced by One Final Audit

A final audit can verify that a calculation exists. It cannot manufacture the calculation.

It can check whether Z_2 was evaluated; it cannot turn a symbolic projection formula into a numerical weak coupling.

It can check whether $\hat{H}$_cl was constructed; it cannot convert a declared closure-Hessian debt into four Yukawa matrices.

It can check whether the global determinant line is trivial; it cannot obtain that trivialisation from local perturbative anomaly cancellation.

It can check whether the triality gap survives the continuum limit; it cannot replace the missing coercivity and quantum-sector proofs.

That is why the programme has eight remaining papers rather than one.

Final Perspective

VERSF is now close to the end of its Standard Model theorem architecture.

That is a substantial achievement in its own right. The gauge group, matter ledger, scalar carrier, Yukawa form, flavour-frame logic, quantum gauge machinery, running architecture and conditional confinement structure are all placed inside one connected programme, with named premises, falsifiers and dependency maps.

But the word **derivation** ultimately carries a numerical and operational burden. A theory that fixes the form of every Standard Model object but the value of none has completed its architecture, not its derivation.

The final phase must therefore calculate the substrate objects, freeze the boundary data, run the theory, match it to physical observables — and publish the result before inspecting whether nature agrees.

The eight remaining papers are not a sign that the programme has failed to converge. They are the final conversion from

a complete derivational architecture

to

a complete, renormalised and auditable Standard Model calculation.

The last paper will then reach one of exactly three honest conclusions:

1. VERSF derives the Standard Model within its declared scope;
2. VERSF derives part of the Standard Model and retains named unresolved parameters or global debts;
3. one or more frozen VERSF calculations disagree with nature and the relevant sector fails.

Any of the three is scientifically clearer than leaving the distinction between architecture and calculation unstated. That is the purpose of the final eight-paper roadmap.