

# How Far Has VERSF Progressed Toward Deriving the Standard Model? June 22<sup>nd</sup> 2026

*An Updated Programme Status Report — Gauge Structure, Matter, Flavour, and the Open Problems*

Keith Taylor · VERSF Theoretical Physics Programme

**What this document is.** This is a status report, not a result paper. Its job is to take honest stock of how far the VERSF programme has travelled toward its central goal — and, just as importantly, how far it still has to go. It is written for the general reader as much as for the physicist. Where it uses the programme's technical vocabulary, it explains it in plain terms first. Nothing here is a new derivation; everything here is an assessment of work done elsewhere in the corpus. This is the second such report, and it is written to be read against the first: where the picture has moved since then, this report says so explicitly rather than quietly overwriting the earlier judgement.

## A Note Before We Begin

It is easy for a programme like this to flatter itself. A document that grades its own work is, by its nature, the least trustworthy document the programme can produce — and an over-confident one is the first thing a sceptical outside reader would use to dismiss the whole effort.

So this report is written against that temptation. It leads with what is genuinely established, it is explicit about what is merely plausible, and it puts the unsolved problems where they belong: in plain sight. A reader should finish it knowing exactly which claims VERSF can stand behind today and which are promissory notes.

One thing cannot be said too plainly: **no internal review substitutes for scrutiny by independent working physicists.** A framework claiming to recover the deep structure of particle physics earns its credibility only when

experts outside the programme try hard to break it and cannot. This report is a map of the territory, not a verdict on it.

## What Has Changed Since the Previous Report

This report succeeds an earlier status report, and it is most useful read against it. Most of the foundational picture is unchanged: the gauge sector, charge quantisation, and the matter skeleton — confinement and fractional charge emerging together from one structure — stand exactly where the previous report left them, and nothing in the recent work weakens them. Where this report restates those sectors, it is confirming the earlier judgement, not revising it.

What has moved is concentrated in three places.

**The generation problem has changed route.** The previous report framed "why three?" as open problem D5, split into a *depth* face and a *capacity* face, and reported that the depth face had been proven independent of the framework's assumptions — reducing it to a single uncomputed number that could come out "three" or "four." This report does not retract that honesty, but it overtakes it: the sharper current lead is the D7 census reading (§7), under which the group has exactly three two-dimensional irreducible representations and a putative fourth folds back into an existing one. D5's "compute one number" framing has not been refuted; it has been superseded by a more specific structural target carrying its own named gates.

**The closure operator has gone from IOU to specification.** The previous report named it the single largest foundational task and the programme's shared chokepoint. That remains true, but the operator is no longer an unnamed promissory note: it is now a *specified* object (§8), with its required properties written down and its construction placed explicitly downstream of the D7 / census decision. The earlier phrase "operator unbuilt" was accurate then; it is too crude now.

**Quantitative mass structure is genuinely new.** The previous report had no dedicated mass-law sector — fermion masses lived inside a single "masses and mixings" section graded mostly on architecture, and the prediction-versus-postdiction discussion had little concrete numerical material to test. This report adds three sectors that did not exist before: the charged-lepton engineering law (§9), the graded quark mass-ratio grid (§10), and the  $\chi$ -column halving programme (§11). This is where the programme has done most of its recent work, and where the scorecard has grown new rows.

A reader comparing the two scorecards will therefore see the structural top of the table essentially unchanged, the generation block relabelled around the D7 route, and a new quantitative-mass block in the middle that the earlier report could not yet write. The overall verdict — a partial, conditional structural account, not yet a derivation — is the same in both. What has changed is that several formerly vague gaps are now named gates with concrete computations attached.

**How the change is marked in this report.** To make that comparison visible at a glance, the scorecard at the end (§17) shows some of its stars in red. ★ **Red stars** are stars the present report holds that the previous one did not — either a grade that has *risen* on a sector graded in both reports (this happens once, for Full operative construction), or, far more often, every star of a sector that is entirely *new* since the previous report (those rows are tagged **NEW**). Black stars are unchanged from the previous report. The intended reading of the scorecard is therefore a table that stays black across its structural top and turns red through its middle: stable foundations, with a new quantitative storey built above them.

## 1. What the Standard Model Leaves Unexplained

The Standard Model of particle physics is the most successful scientific theory ever built. It predicts the outcomes of experiments to a precision that has no rival anywhere in science. And yet, for all its success, it explains remarkably little about why it has the form it does.

It is useful to sort what the Standard Model contains into three kinds of content.

**Gauge structure — the rulebook for the forces.** Nature's forces are described by a specific mathematical object written as  $SU(3) \times SU(2) \times U(1)$ . The Standard Model does not say why nature chose this rulebook rather than some other. It simply writes it down because it works.

**Matter structure — how the particles are organised.** Quarks are permanently locked inside larger particles and never seen alone (confinement). They carry electric charges in thirds. Matter comes in three near-identical copies, called generations, that differ mainly in mass. None of this is fully explained inside the Standard Model; it is observed and recorded.

**Numerical structure — the actual numbers.** The masses of the particles, the strengths of the forces, and the angles describing how particles transform into one another are roughly twenty numbers that the Standard Model cannot predict. They are measured in experiments and typed in by hand.

The animating question of the VERSF programme is simple to state:

How much of this — the rulebook, the organisation, and eventually the numbers — can be made to follow from a few deeper principles, rather than being assumed?

VERSF's wager is that reality is built from irreversible commitment events on a discrete underlying substrate, that time and geometry are not fundamental but emerge from those events, and that the structure of physics is forced by a small set of consistency requirements on how distinguishable things can be compared, transported, refined, and maintained. This report assesses how much of the Standard Model that wager has so far recovered.

## 2. How to Read the Grades in This Report

Every sector below is given a grade. Because the grading scheme is the most abusable part of any status report, it is worth being completely explicit about what the grades mean — and what they do not.

**The grades measure construction, not truth.** A high grade means the structural argument is well developed and internally coherent, judged against the standard of a *conditional derivation*: a derivation that is sound provided its stated assumptions hold. A high grade does not mean the result has been confirmed against experiment, nor that all its assumptions are themselves established. A five-star mechanism resting on an unproven premise is still resting on an unproven premise.

To keep this honest, each sector carries two grades:

- **Mechanism (M)** — how mature and complete the internal construction is, taken on its own terms.
- **Chain to observation (C)** — how complete the whole path is, from first principles all the way to a measured fact about the real world.

These differ whenever a well-built mechanism depends on a premise that is not yet secured. **The chain grade can never exceed the grade of its weakest necessary link.** This inheritance rule is applied strictly below, and it is the single biggest difference between this report and a naive scorecard.

The five-star scale: ★☆☆☆☆ early idea · ★★☆☆☆ partial · ★★★☆☆ substantial but incomplete · ★★★★☆ strong, conditional · ★★★★★ mature conditional construction.

### 3. Gauge Structure — Why Forces Need a Rulebook At All

**The plain version.** Before asking which rulebook nature uses, one can ask why forces obey a rulebook of this mathematical type at all. VERSF's answer is that the moment you insist on comparing things in different places in a consistent way, you are forced to introduce exactly the machinery — connections and transport — that physicists call gauge structure. It is not an arbitrary choice; it is the price of consistent comparison.

From this starting point the programme develops transport connections, the loop structure physicists call holonomy, the emergence of electromagnetism in its Maxwell form, the built-in redundancy of gauge descriptions, and both the simple U(1) phase sector and the more intricate non-Abelian sectors.

This remains one of the most mature parts of the programme. The idea that *some* gauge structure is mandatory, rather than optional, is argued from genuinely deep principles. The later matter and flavour work does not weaken this sector; if anything it makes clearer why transport and comparison sit near the root of the whole programme.

	GRADE	DEPENDS ON
Mechanism	★★★★★	—
Chain to observation	★★★★☆	full Maxwell–transport identification

#### 4. The Rulebook Itself — Why $SU(3) \times SU(2) \times U(1)$ ?

**The plain version.** Having argued that a rulebook is mandatory, the harder task is to explain why nature picked this particular rulebook. Several independent lines of argument now point the same way: U(1) from a compact transport phase, SU(2) from binary assignment and orientation, and SU(3) from three-channel transport participation.

A recurring three-by-two pattern in the framework's closure architecture is argued to be the source of characteristic Standard-Model multiplicities, tracing them to the unique factorisation

$$6 = 3 \times 2$$

within the closure architecture. The newer quark work sharpens this by separating the channel sector  $C$ , with dimension 3, from the fold sector  $G$ , with dimension 2, inside a residual decomposition  $R = C + G$ . That separation is useful because it prevents the programme from spending the same "three" twice: the colour/channel three and the generation or fold structures are different objects.

The gauge-group story is coherent and increasingly well connected to the matter architecture. It is still conditional on the closure architecture being exactly what the programme says it is, and on the relevant identifications being made without selecting them after seeing the answer.

	GRADE	DEPENDS ON
Mechanism	★★★★☆	uniqueness and audit of the closure architecture
Chain to observation	★★★★☆	same

## 5. Charge Comes in Whole Numbers — Why?

**The plain version.** Electric charge always comes in fixed units; you never find a particle with 1.37 units of charge. In the Standard Model this is simply assumed. VERSF argues it follows from topology — specifically that the transport phase lives on a circle, and you can only wind around a circle a whole number of times. Whole numbers of windings give whole units of charge.

The remaining gap is the identification of electromagnetic phase with transport phase itself. If that identification holds, charge quantisation stops

being an input and becomes a consequence of shape. The mechanism is strong; the chain waits on closing that one link.

	GRADE	DEPENDS ON
Mechanism	★★★★☆	compact phase / winding construction
Chain to observation	★★★★☆	full Maxwell–transport identification

## 6. Confinement and Thirds of Charge — The Cleanest Matter Result

**The plain version.** Quarks carry charges in thirds and are never seen on their own. These two facts are usually treated as separate puzzles. VERSF's most convincing matter-sector result is that they are the same puzzle. The whole-number winding is owned by a closure class as a unit; the fractional shares are what individual members register inside it. The very structure that produces fractional shares is what forbids a free fractional particle from ever escaping.

This is important because the framework does not generate fractional charges and then bolt on a separate mechanism to hide them. One structure does both jobs at once. This remains among the most persuasive sectors of the programme.

Recent papers have not displaced this result. They have made its boundary cleaner: quark/lepton participation, colour-channel participation, and closure maintenance now appear as different faces of one larger participation architecture. But the core result — fractional charge and confinement emerging together — still stands as the strongest matter-sector claim.

	GRADE	DEPENDS ON
Mechanism	★★★★★	—
Chain to observation	★★★★★	—

## 7. Matter Comes in Three Copies — No Longer Just "Almost Explained"

**The plain version.** The electron has two heavier relatives, the muon and tau; the same threefold pattern repeats across the matter particles. Why three? Earlier versions of this report said VERSF could generate three-level structures but had not settled why the sequence stops at three. That remains true, but the situation has changed.

The newer work splits the problem into sharper pieces.

First, the framework now separates **assignment, participation, and realization depth** as different axes. Assignment accounts for weak doublet structure, participation accounts for quark/lepton and colour behaviour, and realization depth accounts for generation structure.

Second, the replication work gives a cleaner reason why families can repeat with the same charges: charge reads winding and seat data, while mass can read realization depth and completion density. Charge identity and mass difference are no longer blurred.

Third, the realization operator has been **specified** — which is not the same as **constructing** it. The specification says what the operator must do (preserve closure, yield finite distinguishability, isolate stable modes, be compatible with saturation) but it does not by itself force exactly three modes. The old

phrase "operator unbuilt" is therefore now too crude. The problem is no longer unnamed; the operator's required properties and construction dependencies are explicit.

Fourth, the **D7 route** is a serious new lead. Under the leading reading of the fourteen-generator census as the order of D7, the group has exactly three two-dimensional irreducible representations, and a putative fourth folds back into an existing one. If realization eigenmodes are those 2D irreps, and if the multiplicity-free "acts-on" reading is right, then three-ness and fourth-generation exclusion would follow from the group rather than being assumed. But the *ifs* are load-bearing: the census reading, the multiplicity issue, and the physics identification all remain open.

So the honest update is this: the generation problem has moved from a broad open question to a set of named gates. That is a major improvement, but it is not yet a completed derivation of three observed generations.

RESULT	MECHANISM	CHAIN TO OBSERVATION	DEPENDS ON
Generation architecture	★★★★☆	★★☆☆☆	realization-depth mapping, O1
Replication with charge identity	★★★★☆	★★☆☆☆	charge-blind depth index, realization mapping
Realization operator specification	★★★★☆	★★☆☆☆	actual operator construction
D7 route to exactly three	★★★★☆	★★☆☆☆	D7 census, multiplicity-free carrier, eigenmode identification

RESULT	MECHANISM	CHAIN TO OBSERVATION	DEPENDS ON
Fourth-generation exclusion	★★★★☆	★★☆☆☆	same D7 gates

## 8. The Closure Operator — From Vague Chokepoint to Specified Chokepoint

**The plain version.** Several parts of VERSF now point to the same missing machine: the concrete operator that acts on admissible closure states and whose spectrum gives the stable realization modes. The old status report called this the largest single foundational task. That is still right, but the task is now much more sharply described.

There are two developments to keep apart.

The **closure-state-machine** work shows that some earlier realizability caveats can be discharged in finite models. That is real progress: it shows the programme is not only naming objects but beginning to exhibit state machinery.

The **realization-operator specification** then says something more disciplined: the physically load-bearing realization operator is specified, not built. Its construction is downstream of the fourteen-generator census and the D7 / closure-architecture decision. Until that decision is settled, "closure geometry" is not yet detailed enough to write the operator action.

This is a better situation than the old one. The problem is no longer "we need an operator" in the abstract. It is: settle the closure architecture, choose the correct carrier reading, construct the realization operator, and compute its spectrum. Several gates that used to look separate now converge on this same task.

	GRADE	DEPENDS ON
Mechanism	★★★★☆	D7 / census decision, carrier choice
Chain to observation	★★☆☆☆	explicit operator action and spectrum

## 9. Charged Leptons — The First Concrete Mass Engineering Law

**The plain version.** The Standard Model records that the muon is about 207 times heavier than the electron, and the tau is about 16.8 times heavier than the muon. It does not explain either number. VERSF now has its strongest quantitative mass result here. The previous report had no sector like this one: charged-lepton masses appeared only inside a general "mass architecture," graded on whether the machinery was computable rather than on a specific law. That a concrete ratio law can now be written down at all is itself the advance.

The charged-lepton law has two pieces.

The first is **localization**. A successive generation is more tightly localized, and mass scales like inverse size squared. With the localization exponent  $\kappa = 8/3$ , the electron-to-muon step gives  $e^{(16/3)}$ , which lands within about 0.17% of the observed muon/electron ratio. The caveat is important: the factor of two in the exponent must be grounded by the mass-localization scaling, not smuggled in. Recent work states this dependency plainly.

The second is **saturated participation**. Applying the same localization step again would make the tau far too heavy. The tau needs suppression of about one twelfth. The newer closure-cell work gives that one twelfth a structural

source: the  $W_7$  closure cell has a primitive spectral weight fraction  $2/24 = 1/12$ . The charged-lepton ratio law then becomes

$$m_e : m_\mu : m_\tau = 1 : e^{(16/3)} : (1/12) \cdot e^{(32/3)}.$$

This is no longer merely "choose a residual factor for the tau." It is a convergence: the tau suppression is reached from variational energy collapse *and* from the closure-cell / surviving-participation route. The strongest statement is that VERSF has a structural charged-lepton engineering relation, not yet an unconditional proof.

RESULT	MECHANISM	CHAIN TO OBSERVATION	DEPENDS ON
$\mu/e$ ratio	★★★★☆	★★★★☆	$\kappa = 8/3$ and mass-localization exponent
$\tau/\mu$ ratio	★★★★☆	★★★★☆	1/12 activation and surviving-channel rea
full charged-lepton law	★★★★☆	★★★★☆	uniform $\kappa$ , Role-4 scaling, closure-cell co

## 10. Quark Mass Ratios — From Numerical Scatter to a Graded Grid

**The plain version.** Quarks are harder than charged leptons because they are coloured, confined, and scheme-sensitive. The useful move is therefore not to chase absolute masses first, but to build a grid of dimensionless ratios and grade each cell by what it has actually earned.

The conditional quark mass-ratio grid is a genuine advance because it refuses to treat a row of good-looking numbers as one uniform success. Where the

previous report could only say the mass architecture existed and was "computable, not merely descriptive," this report can show the actual ratios alongside an explicit status for each. Normalised to the down quark, the construction places the six quarks roughly as

$$u : d : s : c : b : t = 0.46 : 1 : 20 : 235 : 1,040 : 61,500,$$

against measured values of roughly

$$0.47 : 1 : 20 : 235 : 1,080 : 64,000.$$

But the important row is not the numerical closeness. It is the *status* row.

CELL	STATUS	HONEST READING
d	anchor	chosen normalization, not a prediction
u/d = 6/13	conditional prediction	depends on form provenance, not just K = 7
s/d ≈ 20	import	supplied by Route 1 / kaon-GMOR route, at resolution limit
c/s ≈ 11.7	match	awaits derivation of the $\chi$ law
b/s ≈ 52	conditional prediction	channel saturation: $3 \times (1/12) = 1/4$ , about 4% tension
t/b ≈ 59	match / by-product	depends on running convention and $\chi$ law

The key new quark result is the **bottom/strange step**. A previous 8/3 colour factor looked attractive, but it relied on a scale-mixed target. Once the same-charge b/s target is made RG-invariant, the required enhancement is close to 3, not 8/3. That 3 is then read as the number of surviving colour transport channels: a colourless saturated closure keeps one route, giving 1/12; a coloured saturated closure keeps one route in each of three colour channels, giving  $3/12 = 1/4$ . This predicts b/s to within about 4%.

The **strange/down step is not solved**. It remains imported, and the paper is right to leave it that way. The residual correction is about the same size as the data uncertainty, and channel counting cannot explain it cleanly. This is not a failure hidden under a fit; it is a deliberately exposed boundary.

	GRADE	DEPENDS ON
Mechanism	★★★★☆	K = 7 audit, channel-count descent, cell provenance
Chain to observation	★★★★☆	s/d import, $\chi$ law, scale conventions

## 11. The $\chi$ Column — Empirical Gate Passed, Structural Gate Still Open

**The plain version.** Inside each generation there is an up-type quark and a down-type quark. Define  $\chi$  by

$$\chi(g) = \ln[ m_{\text{up}}(g) / m_{\text{down}}(g) ].$$

The striking pattern is that the increase in  $\chi$  from generation 1 to 2 is about twice the increase from generation 2 to 3. In symbols,  $q = \Delta\chi_2 / \Delta\chi_1$  is close to 1/2.

This  $\chi$  programme did not appear in the previous report at all; it is one of the clearest examples of the recent shift toward auditable quantitative targets. It has now become one of the most disciplined parts of the quark programme.

The **single-scheme  $\chi$  audit** did the empirical housekeeping. It recomputed the three up/down ratios under consistent MS-bar QCD running with four-loop threshold matching. The result is  $\rho = 0.503 \pm 0.023$ , or honestly  $\rho = 0.50 \pm 0.02$ . That removes the main worry that the apparent halving was a pole/MS-bar bookkeeping artefact. It does not measure exact equality, and it does not prove the mechanism. It says the programme is aiming at the right empirical target.

The **binary-fold paper** then gave the structural side. The two-dimensional G sector does not by itself force a halving — that would be too quick. What it does give is an *exchange-odd line*: if the refinement law commutes with the fold exchange, the odd part evolves by a single scalar  $\lambda$ . That forces a geometric profile with some ratio  $\lambda$ . The remaining task is to derive  $\lambda = 1/2$  from the actual access and closure machinery.

The most recent paper advances this further. It does not claim to prove the physical quark gate halves. It gives a complete **conditional reduction**: if six named conditions hold, exact halving follows algebraically. Those conditions are:

1. reversal-fixed external signatures,
2. image-even closure,
3. a quotient-binary one-step history with no flip primitive, equal multiplicity, and equal stabilizer order,
4. equal scalar weight per primitive class,
5. no branch-cost asymmetry,
6. log-access intertwining.

This is a real upgrade. The problem is no longer "maybe symmetry gives 1/2." The paper explicitly says symmetry alone is insufficient: symmetry fixes the averaging *direction* but not the *weight*. The weight comes from uniform

admissibility plus a two-branch census with equal stabilizer and no cost asymmetry. That is exactly the right place to put the remaining work.

The remaining targets are now concrete: compute  $\text{rank}(U)$ , choose the support route or word-span route, test whether the two fold histories are genuine, audit quotient-binarity and stabilizer order, and prove the log-access intertwining without defining it from the desired result. Passing Gate 1 means the fold sector is real; it does not mean it halves.

RESULT	MECHANISM	CHAIN TO OBSERVATION	DEPENDS ON
empirical $\rho \approx 1/2$	★★★★☆	★★★★☆	input precision, QED/EW refinements
geometric $\chi$ profile	★★★★☆	★★★☆☆	G-sector exchange and rate stability
exact 1/2 conditional theorem	★★★★☆	★★★☆☆	six conditions, especially (3), (5), (6)
full c/s and t/b prediction	★★★☆☆	★★☆☆☆	magnitude $\Delta\chi_1$ and anchor proven

## 12. Left-Handedness — A Mechanism, Not Yet the Result

**The plain version.** The weak force treats left and right differently — one of the deepest asymmetries in physics. A long-standing objection to any geometric approach is that geometry seems to have no way to tell left from right in the first place. VERSF has several routes that produce handed behaviour, and work on chiral transport shows concretely that geometry can select one handedness over another.

That is a genuine advance against the old objection. But honesty requires drawing a hard line here. Showing that geometry can produce a left–right asymmetry in transport is not the same as deriving the specific parity violation of the weak force, which lives in the electroweak coupling itself. The programme has a credible mechanism and a proof of principle; it does not yet have a derivation of weak parity violation.

	GRADE	DEPENDS ON
Mechanism	★★★★☆	—
Chain to observation	★★★★☆	bridge from geometric handedness to electroweak coupling

### 13. The Higgs and the Masses of the Force Carriers

**The plain version.** The Standard Model gives mass to the W and Z force carriers through a fundamental field — the Higgs — that it postulates. VERSF reinterprets the Higgs not as a fundamental ingredient but as the "thickness" mode of a condensate that forms when closure structure settles into its preferred state. From this it derives the condensation itself, the splitting that gives the W and Z their masses, the protection that keeps the photon massless, and an interpretation of the Higgs particle as a ripple in that condensate.

This remains a strong sector. It is less central to the recent batch of papers than the lepton and quark hierarchy work, but the existing status does not need downgrading. The caveat is the same as before: any numerical closeness of the Higgs scale must be audited under the prediction-versus-postdiction standard.

	GRADE	DEPENDS ON
Mechanism	★★★★☆	closure-condensate construction
Chain to observation	★★★★☆	channel-counting and provenance audit

## 14. The Masses and Mixings of the Matter Particles

**The plain version.** The Standard Model's least explained content is its list of fermion masses and the angles describing how particles of different generations turn into one another. VERSF now has several partial mechanisms in this territory.

The mass picture has matured. The old phrase "mass = localization cost" is no longer enough. The newer account is closer to

mass = localization × admissibility maintenance.

Localization gives the shared ladder. Maintenance measures how hard a closure has to work to remain admissible. Complete closures, such as leptons, can pay that internally; partial closures, such as quarks, pay through the confinement bath. This makes sense of why the same localization ladder works cleanly for muon/electron, needs a 1/12 suppression for tau/muon, and has to be modified in the quark sector.

The flavour-mixing picture is also cleaner. CKM and PMNS mixing are now framed as misalignment between completion and transport/readout structures. This explains why quark mixing is small while neutrino mixing is large — the quark hierarchy  $|V_{us}| \gg |V_{cb}| \gg |V_{ub}|$  falling out of the geometry — but the exact matrices are still not computed.

Here the inheritance rule bites hard. Many of these mechanisms are now strong internally, but exact measured masses and exact mixing matrices still require the operative realization and access operators. The newer papers have replaced vague gaps with named gates, but they have not closed all the gates.

RESULT	MECHANISM	CHAIN TO OBSERVATION	DEPENDS ON
fermion mass architecture	★★★★☆	★★★★☆	exact realization/access operators
charged-lepton ratios	★★★★☆	★★★★☆	closure-cell activation and $\kappa$ chain
quark ratio grid	★★★★☆	★★★★☆	cell provenance and $\chi$ gates
$\chi$ profile	★★★★☆	★★★★☆	six-condition audit
CKM / PMNS framing	★★★★☆	★★☆☆☆	exact matrices, O1
CP violation	★★★★☆	★★☆☆☆	exact phase and mapping

## 15. Prediction Versus Postdiction — The Test That Matters

This section remains essential.

There is a world of difference between explaining a number you already knew and predicting a number before it is measured or before the framework has been tuned around it. Only the second carries decisive weight, because the first is always vulnerable to the suspicion — fair or not — that the framework was quietly adjusted to fit.

The recent papers improve the programme's discipline here. Several results are explicitly labelled as imports, matches, conditional predictions, or by-products. The quark grid is especially useful because it does not let a row of good-looking percentages hide the fact that some entries are borrowed and some are genuinely structural.

The strongest current numerical candidates are:

- the muon/electron ratio from  $e^{(16/3)}$ , conditional on the  $CP^2 / \kappa$  and mass-localization chain;
- the tau suppression from the  $W_7$  closure-cell unit  $1/12$ ;
- the bottom/strange prediction from  $3 \times (1/12) = 1/4$ ;
- the  $\chi$ -profile target  $\varrho \approx 1/2$  surviving a single-scheme audit;
- the fine-structure ( $\alpha^{-1} \approx 137.034$ ) and Higgs-scale postdictions from earlier work.

But a sceptical reader is still entitled to ask, for each one: was the form fixed before the measured value was known, or selected because it hit the value? The programme should not argue with that question. It should answer it with provenance audits.

The status report's standard is therefore unchanged: a true prediction must follow from the framework without using the answer as a target, and it must remain testable in a way that could have failed.

## 16. The Open Problems, Stated Plainly

These are the load-bearing unknowns. They are listed here because the programme's honesty depends on them being visible.

**G-OP — Construct the operative realization operator.** The operator has been specified and constrained, but the physical action on admissible states is not yet built. The D7 / census result is upstream of that construction.

**D7 / census discrimination.** The leading D7 reading gives a powerful route to exactly three two-dimensional realization modes, but the carrier reading and multiplicity issue must be settled. The "14" must function as group order acting on the carrier, not simply as a 14-dimensional regular representation, unless extra postulates are added. This is the successor to the previous report's open problem **D5** — the depth/capacity split on "why three?": the underlying question is the same, but it is now pursued through the group structure rather than through an isolated, separately computed gap number.

**O1 — Are the framework's internal levels the real generations?** This remains consequential. Some recent work narrows the problem, but exact flavour claims still need the mapping from internal realization objects to observed particles.

**G-UD(form).**  $K = 7$  and the dimension counts may be independently grounded, but selected algebraic forms such as  $6/13 = (K - 1)/(2K - 1)$  need provenance audits. Establishing the integer is not the same as establishing every expression built from it.

**G-LOC2.** The factor  $e^{(16/3)}$  depends on both  $\kappa = 8/3$  and the mass-localization exponent. The factor of two must remain explicitly grounded.

**G-ACT / participation activation.** The closure-cell unit  $1/12$  is strong, but the rule determining which sectors and which steps wear the unit is not fully derived.

**G-CHI structural audit.** The latest  $\chi$  paper reduces exact halving to six conditions. The most exposed are quotient-binarity, equal stabilizer order, no branch-cost asymmetry, and log-access intertwining. These now need computation against the actual A, B, C operators and access maps.

**Exact spectrum.** The architecture for masses exists; the complete fermion spectrum has not been computed.

**Exact CKM and PMNS values.** The mixing mechanisms exist; the precise matrices do not yet follow.

**Prediction provenance.** Several numerical agreements are impressive, but the programme must prove which are genuine predictions and which are postdictions or selected forms.

## 17. The Scorecard

Read against the two-grade scheme of §2. M = maturity of the internal mechanism. C = completeness of the whole chain to a measured fact, capped by its weakest necessary link.

**Reading this scorecard against the previous report.** ★ **Red stars** are stars the previous report did not have. On a sector graded in *both* reports, a red star is a grade that has risen (this happens once — Full operative construction). On a sector tagged **NEW**, the whole sector is new since the previous report, so every one of its stars is new — and this is where most of the recent work shows up. Unmarked grades are unchanged; the structural foundations in particular hold exactly.

SECTOR	M	C	CAPPED BY
Gauge necessity	★★★★★	★★★★★	Maxwell–transport identification
U(1) emergence	★★★★★	★★★★★	full electromagnetic identification
SU(2) emergence	★★★★★	★★★★★	assignment/orientation identification

SECTOR	M	C	CAPPED BY
SU(3) emergence	★★★★★	★★★★★	transport-channel derivation
Gauge-group structure	★★★★★	★★★★★	closure-factorisation audit
Charge quantization	★★★★★	★★★★★	Maxwell–transport identification
Confinement	★★★★★	★★★★★	—
Fractional charge	★★★★★	★★★★★	—
Matter skeleton	★★★★★	★★★★★	—
Generation architecture	★★★★★	★★★☆☆	D7 / census, O1
Replication theorem <small>NEW</small>	★★★★★	★★★★★	realization mapping
Realization operator specification <small>NEW</small>	★★★★★	★★★☆☆	operator action not built

SECTOR	M	C	CAPPED BY
D7 route to three NEW	★★★★★	★★★☆☆	multiplicity-free carrier and eigenmode
Closure-cell suppression unit 1/12 NEW	★★★★★	★★★★★	activation and coupling
Charged-lepton ratio law NEW	★★★★★	★★★★★	$\kappa$ chain, scaling, 1/12 activation
Quark mass-ratio grid NEW	★★★★★	★★★★★	cell provenance, s/d import, $\chi$ law
b/s channel prediction NEW	★★★★★	★★★★★	dim C = 3 descent, step attachment
$\chi$ empirical audit NEW	★★★★★	★★★★★	input precision, QED/EW refinement
$\chi$ -halving conditional theorem NEW	★★★★★	★★★★★	six-condition structural audit
Chirality	★★★★★	★★★★★	electroweak bridge
Electroweak symmetry breaking	★★★★★	★★★★★	closure-condensate construction

SECTOR	M	C	CAPPED BY
Higgs emergence	★★★★★	★★★★★	provenance of channel-counting argu
Fermion mass architecture	★★★★★	★★★★★	exact spectrum
CKM / PMNS framing	★★★★★	★★★★★	exact matrices, O1
CP violation	★★★★★	★★★★★	exact phase and mapping
Exact masses	★★★☆☆	★★★☆☆	open
Exact flavour spectrum	★★★☆☆	★★★☆☆	open
Full operative construction	★☆☆☆☆	★☆☆☆☆	explicit realization/access operators

The pattern is worth reading off directly. The strongest sectors are structural: gauge necessity, confinement, fractional charge, and the closure-cell suppression unit. The most active frontier is now quantitative mass structure: charged leptons, the quark grid, and the  $\chi$  column. The weakest sectors are still the exact spectrum and exact observed mixing values, because they depend on operators that are specified or partially constructed but not yet fully instantiated.

The colouring tells an honest story, and it is worth reading the two kinds of red apart. Only one sector graded in *both* reports has actually risen in grade — Full operative construction, up one star on each of M and C, as the closure operator moved from an unnamed task to a specified object. All the other red sits on **NEW** rows: whole sectors — the closure-cell suppression unit, the charged-lepton law, the quark mass-ratio grid, the b/s prediction, the two  $\chi$  rows, and the new generation routes — that the earlier report had no place for at all. The structural sectors at the top hold their grades exactly. So the picture is not stasis: it is a programme that has left its foundations untouched while building an entire new quantitative storey above them.

Set beside the previous report's scorecard, the shape of this one has changed in a specific way. The top structural block is unchanged, grade for grade. The generation block has been relabelled — the previous report's single "generation architecture" and "census closure (depth/capacity)" rows are now organised around the D7 route, the replication theorem, and the realization-operator specification. And the middle of the table carries an entirely new quantitative-mass block — the closure-cell suppression unit, the charged-lepton ratio law, the quark mass-ratio grid, the b/s channel prediction, and the two  $\chi$  rows — none of which had a row to stand in before.

## 18. Where the Programme Stands

VERSF has travelled well beyond foundational speculation. It now offers coherent structural accounts — conditional, but real — of why forces need a rulebook, why that rulebook looks Standard-Model-shaped, why charge comes in whole numbers, why quarks confine and carry thirds, how electroweak symmetry can emerge, how matter families are organised, and why some particle-mass ratios may have structural origins rather than being arbitrary inputs.

The strongest ground remains the matter skeleton: confinement and fractional charge emerging together from one structure is the kind of result that is hard to wave away. The strongest new quantitative ground is the charged-lepton law and the closure-cell suppression unit. The quark sector has also improved

sharply: it is no longer a collection of suggestive ratios but a graded grid with named imports, predictions, matches, and by-products.

The latest  $\chi$  work is a good example of the programme maturing. It does not claim victory. It says the empirical target 1/2 survives audit; it says symmetry alone is insufficient; it reduces exact halving to six explicit structural conditions; and it names the remaining computations. That is how a serious programme should behave.

The single most accurate sentence is therefore careful about its verbs:

VERSF is a partial, conditional structural account of the Standard Model. It has replaced a meaningful fraction of the Standard Model's brute facts with structural arguments, and it has turned several formerly vague numerical hopes into auditable theorem programmes. But it has not yet derived the Standard Model. The decisive remaining work is to instantiate the operators, close the generation and access-map gates, compute the spectrum, and prove which numerical agreements are genuine predictions rather than selected postdictions.

The work ahead is no longer mainly about explaining why Standard-Model-like structure appears. It is about building and auditing the machinery that would make the explanation quantitative: the realization operator, the closure spectrum, the quark  $\chi$  access map, the activation rule for suppression, and the exact route from internal realization states to observed particles. That is a harder standard — and a better one.

★ **Red stars** are stars the previous report did not have — a risen grade on a sector graded in both reports, or (on **NEW** rows) an entire sector added since then. Most of the movement is in the NEW rows.