From Facts to Forces

How VERSF Is Trying to Derive the Standard Model from First Principles

A general-reader account for people with no background in particle physics

Keith Taylor

VERSF Theoretical Physics Programme · Standard Model Closure Series

General-reader synthesis · August 2026

How to read this paper. No physics knowledge is assumed. Technical terms are introduced only after the everyday idea has been explained. Where VERSF has a promising result but has not yet completed the physical derivation, the paper says so. A close number is not treated as proof merely because it is close.

Modern particle physics has an unusual problem.

We have an extraordinarily successful theory called the Standard Model. It tells us what the smallest known particles do, how the forces of nature act on them, how particles can change from one type into another, and how the Higgs field is involved in giving elementary particles mass.

Its predictions have been tested to astonishing precision. Some of them agree with experiment to more decimal places than any other theory in the history of science.

Yet the theory contains a collection of numbers that it does not itself explain. Scientists measure those numbers in the laboratory and then feed them into the theory by hand. They include the strengths of the forces, the relative masses of the elementary particles, the properties of the Higgs field, and the numbers describing how different types of particle transform into one another. Once those values are supplied, the Standard Model tells us with remarkable accuracy what happens next. What it generally does not tell us is why nature chose those particular values in the first place.

Here is a simple analogy. Imagine an extremely accurate musical instrument that can predict every sound it will make once you tell it the tension of each string and the position of every tuning peg. The Standard Model is far more sophisticated than that, but the puzzle is similar: why are the tuning settings of nature what they are?

VERSF is an attempt to go one level deeper.

Instead of taking all of those measured numbers as fundamental, it asks whether at least some of them could be consequences of a smaller and more basic set of principles. The ambition is not merely to reproduce the Standard Model. It is to explain why something resembling the Standard Model should appear at all, and ultimately why its numbers have the values we observe.

This paper explains how far that programme has progressed and, just as importantly, how the derivation has actually worked: what has been calculated, what has succeeded, what has failed, and what the programme refuses to claim.

Before any of that can make sense, we need to meet the cast.

---

1. Meet the particles: what the world is actually made of

Everything you have ever touched is made of atoms.

An atom is mostly empty. At its centre sits a tiny, dense core called the nucleus. Around that core, electrons are arranged.

The **electron** is, as far as anyone can tell, a genuinely elementary particle: it is not made of anything smaller. It carries electric charge. Electrons are responsible for chemistry, for electricity, and for the fact that your hand does not pass through the table. When atoms share or exchange electrons, molecules form. When electrons flow along a wire, you get an electric current.

The nucleus, in contrast, is not elementary. It is made of protons and neutrons packed tightly together. And protons and neutrons are themselves not elementary either. Each is made of three smaller particles called **quarks**.

Ordinary matter needs only two kinds of quark, with the rather plain names **up** and **down**. A proton is two up quarks and one down quark. A neutron is two down quarks and one up quark. That is the whole recipe.

So the everyday world, in its entirety, is built from just three elementary matter particles:

- the electron,
- the up quark,
- the down quark.

There is a fourth member of the family: the **neutrino**. Neutrinos are extraordinarily light, electrically neutral particles produced in certain radioactive decays and in enormous numbers by the Sun. They interact so feebly with ordinary matter that trillions of them pass through your body every second without leaving a trace. They matter to this story later.

Physicists group these particles into two broad kinds. The quarks are the ingredients of protons and neutrons. The electron and the neutrino belong to a group called leptons, particles that do not feel the strong force. You will not need the word again, but you may meet it elsewhere.

These particles are held together, pushed apart, and transformed by a small number of forces:

- The **electromagnetic force** binds electrons to nuclei, makes chemistry possible, and is carried by light.
- The **strong force** glues quarks together inside protons and neutrons, and ultimately helps hold atomic nuclei together. It is the most powerful force known, but it operates only over distances smaller than a nucleus.
- The **weak force** does something the others cannot: it changes one type of particle into another. It is why certain atomic nuclei are radioactive, and it is one of the reasons the Sun shines.
- Finally there is the **Higgs field**, an invisible background that fills everything and is involved in giving elementary particles their mass. We will return to it in detail, because it turns out to be one of the most developed parts of the VERSF story.

Gravity, incidentally, is not part of the Standard Model at all. That is a separate story.

If nature had stopped there, physics would already be strange enough. It did not stop there.

2. The photocopies nobody ordered: the three families

Here is one of the deepest unsolved puzzles in physics, and it can be stated without a single equation.

Nature contains extra copies of the matter particles.

The electron has a sibling called the **muon**. The muon behaves in almost every respect exactly like an electron: same electric charge, same interactions, same character. There is only one important difference. The muon has about two hundred times the mass of the electron. Being unstable, it survives for about two millionths of a second before decaying. When the muon was discovered in the 1930s, it was so unexpected that one famous physicist reacted with the words, “Who ordered that?”

Then it happened again. The electron and the muon have a third, still more massive sibling: the **tau**.

The quarks do the same thing. The up quark has two more massive copies, called **charm** and **top**. The down quark has two more massive copies, called **strange** and **bottom**. Even the ghostly neutrino comes in three types.

Physicists arrange this into three columns, called the three **families** of matter. You will also see them called the three **generations**: the two words mean exactly the same thing. Popular accounts tend to say families, while technical writing more often says generations. This paper says families throughout. The three are:

- **Family 1:** electron, its neutrino, up quark, down quark. Everything around you.
- **Family 2:** muon, its neutrino, charm quark, strange quark. More massive copies.
- **Family 3:** tau, its neutrino, top quark, bottom quark. More massive still.

The second and third families are like photocopies of the first, identical in every property except mass, and unstable: left to themselves, their particles rapidly decay back towards the familiar first family. They can be produced in energetic processes, in particle accelerators and in cosmic rays, but they do not stay.

The word generation, incidentally, carries a helpful intuition of its own. Each generation is a complete new edition of the entire four-member set, the way successive generations of a product replace the whole line rather than one item in it. So the electron, muon and tau are not themselves a generation: they are the same member, the charged lepton, appearing once in each of the three successive generations.

The Standard Model describes all of this perfectly. It does not explain any of it. It does not say why there are three families rather than one or seventeen. It does not say why the copies have the masses they do. It does not say why the families can partially transform into one another, in the very particular pattern experiment reveals.

3. The mysteries we are trying to explain

VERSF is not trying merely to reproduce familiar equations in a different notation. The ambition is to explain why a theory resembling the Standard Model has the particular structure and numerical settings that it does.

With the cast now introduced, the main mysteries can be stated plainly:

- Why are there stable elementary particles at all, instead of only temporary disturbances?
- Why do the basic matter particles come in three related families?
- Why do the second and third families have so much more mass than the ordinary electron and light quarks?
- Why do the electromagnetic, weak and strong interactions have their particular strengths?
- Why can one member of a quark family transform into a mixture of the others, and why does that mixing have its peculiar pattern?
- Why does nature show a small but crucial difference between matter and antimatter?
- Why is a particular matter-antimatter quantity in the strong interaction extraordinarily close to zero?
- Why do neutrinos have such tiny masses, and why do they change between observable types?
- Can all of these results come from one underlying physical construction, rather than from separate fitted explanations?

Those questions are related. A satisfying deeper theory should not answer each one by adding a new adjustable number. Ideally, the same underlying rules that create persistent matter, record history and permit interactions should also constrain the numerical values that the Standard Model presently takes from experiment.

And the real test is not whether a clever formula can land near a measured number. It is whether the deeper theory was forced to produce that number before the answer was opened.

To attempt this, VERSF starts from a place that may seem surprisingly far away from particles altogether.

4. Starting before particles

Most theories of particle physics begin by assuming that things such as particles, fields and a geometric arena already exist.

VERSF tries to start earlier.

Its basic picture is that reality progresses through a succession of **irreversible facts**.

Think of a fact simply as something that has become settled. Before a fact is produced, several alternatives may still be possible. When one alternative becomes physically committed, the world has acquired something that was not there before: a piece of history that cannot be treated as though it never happened.

The next fact then forms in the context of the facts that already exist.

So in VERSF, time is not initially introduced as a clock ticking against an empty background. The ordering

fact 1 → fact 2 → fact 3 → fact 4

is itself what gives physical events their before-and-after structure. The technical VERSF work describes this as a handoff: one fact becomes fixed while the next begins to form.

This apparently philosophical starting point quickly becomes mathematical, because several questions must be answered precisely:

- How does one possibility become the realised fact?
- How is that fact remembered?
- How is a fresh set of possibilities created without deleting what has already happened?
- And can the accumulated history affect what happens next?

The recent VERSF programme has spent a great deal of effort answering exactly these questions.

5. A fact must leave a record, and the present must reopen

Suppose you write something in a notebook and then turn the page.

Two failures are possible, and they fail in opposite directions. If every new sentence erased the previous page, there would be no history. But if the notebook froze permanently after the first sentence, nothing new could ever happen. The old page must remain written, and there must also be a fresh page on which something new can occur.

VERSF has developed exactly this distinction between the **persistent record** and the **active present**.

The latest calculations show, within the current architecture, that both requirements can be met at once: a committed fact can be retained while the active working system is reopened for the next event. This is better described as retention plus renewal than as reality making a new copy of itself every moment. The record stays in the growing history while the active part of the system renews.

This distinction has turned out to matter enormously, because of a subtlety that is easy to miss.

Keeping the past is not automatically the same as the past influencing the future.

A library can preserve every book ever written without the books themselves reaching out and changing what someone writes tomorrow. Whether the books matter depends entirely on whether anyone is reading them.

VERSF has turned that observation into two precise results, one for each side of the question.

The first result concerns a version of the dynamics in which the stored record and the active system are kept disconnected, each evolving on its own. For that version, VERSF proved that the history is perfectly preserved and yet influences nothing at all: an unread library. The value of this result lies in what it rules out. Storage alone can never produce influence. So if the past is to shape the future, there must be a genuine physical connection through which the stored information reaches the active system. The theorem tells you the wiring is necessary.

The second result concerns dynamics in which such a connection is present. There, VERSF has shown that different histories genuinely lead to different probabilities for what happens next, and that a state-preserving renewal carries the influence forward from one fact to the next instead of erasing it. The library, once read, changes what gets written.

The two results fit together without tension: no connection, no influence, guaranteed; connection, influence, demonstrated. The remaining problem is sharply defined: show which renewal and which history-to-present connection are *forced* by the primitive dynamics, rather than merely allowed by them.

This is a good example of how the programme has developed. A question that once sounded hopelessly vague ("how can the past affect the present?") has become a specific mathematical question about particular physical maps.

6. From continuing patterns to matter

Once facts can persist and history can grow, a second question appears:

why are there enduring *things* at all?

In VERSF, a particle of matter is not pictured as a tiny hard bead.

A useful first analogy is a whirlpool. The water making up a whirlpool is constantly moving. Individual water molecules enter and leave, yet the whirlpool preserves its overall identity. It is a pattern that persists, made of nothing but the medium it lives in.

VERSF proposes something similar for matter. A particle is a persistent pattern in the underlying structure: a pattern capable of surviving from one fact to the next.

There is an important difference from an ordinary whirlpool, however. A whirlpool dies when the water stops moving. The persistence in VERSF is intended to be **topological**: more like a knot tied in a rope than a swirl in water. You cannot smooth a knot away by gently deforming the rope. You would have to cut the rope to remove it. And a knot persists because of how the rope is arranged, not because it is made from special rope. The identity of a particle, in this picture, is associated with how the underlying structure is wound or closed, and that winding is what protects it.

This leads to a clean division of labour:

- the succession of facts supplies the ordering we experience as time;
- persistent closure supplies the identity of matter;
- the Higgs-related background supplies a common elementary mass scale;
- the detailed structure of a particular particle determines how strongly it responds to that scale.

Keep that last pair in mind. It becomes important immediately.

7. What the Higgs means in this picture

The Higgs is often popularly described as "the thing that gives particles mass". That is not entirely wrong, but it invites the wrong mental image, as though the Higgs were treacle that particles wade through.

VERSF uses a different analogy: a stretched surface.

Imagine a large sheet of elastic material held at a stable, uniform tension.

Two quite different things can exist in that sheet at the same time. A persistent pattern can live in it, rather like a knot or a self-holding whirl, made of the sheet itself. And a disturbance can travel across it as a ripple.

VERSF associates matter with the persistent pattern, and a Higgs excitation with the travelling ripple.

Notice what this picture does *not* say. The Higgs background is not continually rebuilding the particle, moment by moment. The particle already has its identity from its closure structure, the knot. What the Higgs background supplies is a common elementary mass scale: a single shared tension against which different particle structures respond differently. The knot provides the identity; the tension of the sheet sets the scale; the particular shape of each knot decides how strongly it feels that tension.

The programme has now identified a particular common motion of the relevant Higgs-related degrees of freedom. Picture four points on the stretched sheet. Lift one of them and the sheet between them is strained. Lift all four together, by the same amount, and nothing is strained at all: the entire compatible system moves as one. That common, strain-free motion is the current VERSF Higgs direction.

This common motion is unusual in a useful way. The part of the theory that measures mismatch cannot see it, because when everything moves together nothing is mismatched. Yet the stored energy does see it: there is a preferred stable background, and moving the whole structure away from that background costs energy. So the background resists displacement, and disturbances of it can propagate from one fact to the next.

This is why the Higgs sector is currently among the most developed areas of the VERSF Standard Model programme, and why it produces the programme's first striking number.

8. The first striking number: 32/63

One of the strongest numerical results to date concerns the Higgs.

VERSF produces the exact ratio

32/63 = 0.5079365…

for a dimensionless Higgs-related quantity: roughly speaking, the stiffness associated with a Higgs disturbance compared with the size of the stable Higgs background.

Only after that value is fixed is it compared with the corresponding number inferred from the measured Higgs mass and the experimentally known electroweak scale.

The difference is about 0.05 per cent.

The significance is not simply that 32/63 is close. It is that 32/63 is a simple structural number, a ratio of two small whole numbers, obtained *before* consulting the measured Higgs value. VERSF deliberately distinguishes that from the ordinary practice of fitting a decimal to experiment. A fitted number can always be made to agree; a number the structure hands you in advance cannot.

But there is an equally important qualification, and the programme states it openly.

It has not yet been shown that this ratio is the final value returned by the completely physically normalised theory. Nor has the programme completely settled how a Higgs disturbance propagates through physical space. Two microscopic descriptions currently lead to different low-energy behaviours: one gives the ordinary spatial propagation expected of a wave, while the other, after an auxiliary current is allowed to relax fully, begins only at a higher order. The primitive theory must decide which mechanism is actually present.

The next Higgs problem is therefore extremely specific: does the primitive theory contain an independent, reversible neighbour-to-neighbour stiffness or not? If it does, the theory must return its size and physical normalisation. If it does not, the other route must supply the physical spatial propagation. Until that fork is resolved, the Higgs mechanism is graded “structural Stage 3” rather than completed Stage 3 (the grading system is explained in Section 18).

So the Higgs result is strong, but not finished.

That combination, an impressive calculation together with an explicitly unresolved gate, is typical of the present state of VERSF. You will see it again.

9. From closure to electromagnetism

Another particularly interesting result concerns the strength of electromagnetism.

Recall from Section 1 that the electromagnetic force is what binds electrons to atomic nuclei, allows chemistry to exist, and underlies light. In the Standard Model its strength is represented by a famous measured number. Physicists usually quote it upside down, as roughly **137**: the larger this number, the weaker the force. Its precise measured value is about 137.036, and nobody has ever been able to say why.

VERSF asks whether that number can arise from the structure of fact formation itself.

Two pieces emerge from the current construction.

First, the basic seven-part closure structure at the heart of the framework gives **128 unresolved possibilities** before a fact has distinguished among them. Under the stated principle that there is no prior physical information favouring one unresolved possibility over another, the fully closed possibility carries a weight of exactly

1/128.

Second, a separate interface calculation identifies **fifteen** relevant directions. Fourteen are allowed to transmit once the physical interface behaves in the required way, while one global direction does not. This gives a factor of

14/15.

Combining them:

1/128 × 14/15 = 7/960.

Turned upside down, in the conventional way,

960/7 = 137.142857…

The measured value is about 137.036, so the difference is roughly 0.08 per cent.

Again, the important word is *conditional*.

VERSF has not yet completed the physical bridge from this structural probability to the electromagnetic interaction measured in laboratory physics. Two bridges remain, and both are explicitly identified. The microscopic theory must prove that the relevant physical interface really does transmit exactly those fourteen directions. And the resulting structural quantity must then be connected through the standard large-scale description of electromagnetism, and carried to the scale at which the laboratory value is quoted.

So the present claim is not "VERSF has derived the fine-structure constant."

It is: VERSF has produced a simple, target-independent structural number strikingly close to the electromagnetic strength, and the remaining physical work needed to test the connection is explicitly known.

10. Why the photocopies have different masses

Return now to the three families of Section 2, and to the most obvious question they raise.

The muon behaves in essentially every way like an electron. Why should it have about two hundred and seven times the electron's mass? The Standard Model records the two masses; it does not derive the relationship between them from any deeper principle.

VERSF has explored whether the cost of localisation and persistence across successive structural levels, successive depths of the knot, so to speak, naturally generates those ratios.

One current result returns the value

207.127…

for the ratio of the muon mass to the electron mass. The number is not fitted: it arises from a simple exact expression built out of the framework's own structural quantities (a fixed mathematical growth constant raised to the power sixteen-thirds).

The measured ratio is approximately 206.768, a difference of about 0.17 per cent.

A further step of the same construction, intended to reach the third family's tau, gives a result about 2.6 per cent from the observed ratio.

These results are not all equally strong, and the programme keeps the distinction visible. The muon-to-electron relation is deliberately described as a conditional in-domain result rather than a clean independent prediction, because part of its justification lives inside the very domain being described. The exact physical mechanism behind the final tau step is also still incomplete.

This is an honest middle category: numerical proximity that is encouraging, but whose provenance is not yet as strong as that of a completely independent prediction. The programme labels it as such rather than promoting it. Not every close number is granted the same evidential status, and that is intentional.

11. A harder problem: what it means for the families to mix

Now for the hardest part of the story, and one of the most interesting. It is one of the most confusing Standard Model ideas for a non-physicist, so it is worth building from scratch.

Recall that the weak force can change one type of particle into another. In particular, it can turn one kind of quark into a different kind. When it does so, something odd happens: the transformation does not respect the tidy three-family bookkeeping. A quark from one family can turn, with small probability, into a quark from a *different* family.

Here is an analogy that the VERSF flavour papers themselves use.

Imagine three knots tied in the same kind of rope. Each knot is stable, and each has a slightly different structure. These are the three quark families.

Now imagine a machine that can transform one knot into another. The crucial point is that the patterns naturally stable in the rope and the patterns naturally produced by the machine need not line up perfectly. What comes out of the machine may be mostly knot number two, but with a small admixture of knot one, and a still smaller admixture of knot three.

The table describing those overlaps, how much of each stable knot appears in each transformed one, is what physicists call the **CKM matrix**. If the two sets of patterns lined up perfectly, the table would be trivial and the families would not meaningfully mix. The puzzle is why the real table has exactly the lopsided pattern that experiments reveal.

Buried in that table is one number more precious than the rest: a measure of the slight difference between how matter and antimatter behave, one of the ingredients thought necessary for a universe containing any matter at all.

Experiment knows this table extremely well. The Standard Model does not explain why its entries form their particular pattern.

VERSF tries to derive the pattern from the same geometry that distinguishes the three families in the first place: the mixing table, in this picture, is not an extra ingredient but a shadow cast by the family structure itself.

12. The first blind quark calculation, and why its failure matters

The first VERSF attempt at the mixing table is especially important, for a reason that has nothing to do with success.

It was run as a sealed procedure. Two leading scales were openly calibrated from experiment and declared as calibrated inputs; the remaining comparison quantities were withheld from the calculation. It produced a definite mixing table.

Some entries came out close, but those close entries were heavily influenced by the declared calibration, so the programme does not count them as predictions. What matters is what happened on the quantities the calibration did *not* fix. There the calculation failed badly: one important entry came out about 31 per cent too low, and the matter-antimatter measure described above came out about 41 per cent too low.

VERSF retains that failed matrix as the official controlling result. It is not hidden, softened or replaced.

That is deliberate. A theory learns far more from a blind failure than from a successful formula chosen after the answer is known. And the failure matters all the more because of what came next.

13. How the quark correction emerged

The later calculation did not begin by asking, “what correction would make the table fit?” That road leads nowhere interesting: with enough freedom, anything can be made to fit anything.

Instead it asked a structural question: what kinds of correction are even *allowed* by the underlying sixfold structure, once it is reduced to three particle families?

The answer turned out to be unexpectedly rigid.

A direct, simple correction is forbidden outright: the mathematics of how the sixfold structure can act on three families rules out the obvious linear route. The first correction the symmetry permits is of a specific quadratic form, and that form is unique under the stated symmetry. Moreover, it treats the three families democratically: no family is singled out.

A calculation on the smallest honest physical model then did something stronger than merely allowing that threefold symmetry. It actually returned it.

The work has since been carried into the much larger charged-particle structure, a far bigger calculation, where the same family direction remains extremely stable across two independent levels of calculation, with the response changing only very slightly. The correction is no longer a toy-model curiosity.

One subtlety deserves mention because of how the programme handled it. One calculation still carried the overall size of the correction as a free quantity, while another conditionally fixed that size to a definite value. These are not necessarily in contradiction: if both calculations are using the same physical family generator with the same normalisation, the free quantity in the larger calculation is simply a slot waiting to be filled from upstream. But “not necessarily contradictory” is not “proved identical”, and the programme says so.

The sector has therefore progressed from “perhaps there is a correction” to a much more specific claim: there is a unique lowest-order correction geometry, and an executed physical carrier capable of carrying it. The remaining problem is to prove that the calculations are genuinely descriptions of the same physical quantity with the same normalisation.

And, in keeping with the programme’s habits, the current papers have stated *in advance* what that common calculation should return numerically. Agreement would strengthen the bridge. Disagreement would falsify at least one of the bridge assumptions, and is not allowed to be repaired afterwards by rescaling the answer.

14. A remarkable candidate angle, and why VERSF refuses to use it yet

The quark-mixing correction also has an orientation, an angle, which is related to the matter-antimatter difference mentioned above.

Three independently developed VERSF constructions now restrict the possible orientation in different ways. One allows a pair of possible angles. A second allows a set of three. A third, based on carrying the structure around a closed loop of transformations, gives a further pair.

Remarkably, the only angle common to all three lists is **150 degrees**.

If the 150-degree branch is used, the agreement with the measured quark-mixing data improves dramatically. The standard diagnostic measure of disagreement falls from roughly 290 for the current blind baseline to around 5 for the favourable branch. In this diagnostic, smaller is better, and the improvement is enormous: the worst baseline failure is almost completely repaired, and the matter-antimatter measure improves strongly as well.

It would be tempting to declare victory.

VERSF deliberately does not.

Why? Because choosing 150 degrees *because it gives the better answer* would silently convert the calculation into a fit. The physical theory has not yet proved that the choices made inside the three separate constructions are all the same physical selection. The primitive theory must independently select the orientation. Until it does, the poorly fitting blind calculation remains the official result, and the far better 150-degree calculation remains a quarantined diagnostic: recorded, tested, and pointedly not used.

So the current position is unusual but scientifically clean: 150 degrees is the unique convergent candidate from three differently constructed restrictions, and it gives a dramatically better result, but it is not yet a prediction.

This may be the clearest example of the standard VERSF has imposed on itself:

a beautiful answer is not evidence if the measured answer helped choose it.

15. The strong force: progress and restraint

The strong force, remember, is the interaction that binds quarks into protons and neutrons and ultimately helps hold atomic nuclei together.

The strong force has an awkward property that makes careless comparisons especially dangerous: its measured strength changes with energy. Two perfectly correct numbers can disagree wildly simply because they are quoted at different scales, or in different calculation conventions. Comparing them directly is meaningless.

That is not a hypothetical danger. The VERSF programme previously had a conditional strong-force value that was being compared with the standard low-energy benchmark, and the audit found that the two numbers referred to different scales and conventions. The comparison was explicitly withdrawn rather than left standing as a misleading percentage agreement.

The programme has nevertheless made genuine progress. A history-weighted microscopic calculation produces the precursor value

1.701045…

from upstream VERSF ingredients, without using the measured strength of the strong force at any point.

The current papers explicitly refuse to call this the physical strong-force coupling, because another internal projection and the complete physical reduction are still missing. It is a precursor: a genuine piece of the eventual answer, produced by the theory, that has not yet been carried all the way to the quantity experiment measures.

The correct statement is therefore neither "VERSF predicts the strong coupling" nor "VERSF has no strong-sector number". It has a genuine source-side precursor and a well-defined route to the physical quantity. A precursor honestly labelled is worth more to the programme than a coupling prematurely claimed.

16. One of the strongest structural results: the knob that may not exist

The strong force contains another famous puzzle.

The mathematics of the Standard Model appears to allow a particular quantity, describing a possible difference between matter and antimatter in strong interactions, to take a broad range of values. It is, in effect, a knob the theory is free to turn.

Experiment says the knob is set extraordinarily close to zero. Not roughly small. Astonishingly, unnaturally close to zero.

Why would nature tune a knob so precisely?

Within the current retained VERSF construction, that quantity is not adjusted to be tiny. It comes out **exactly zero**, because the relevant completed structure contains no freedom to choose a non-zero value. There is nothing to turn.

No measurement of the quantity, and no experimental bound on it, was used to select this result.

The result remains conditional, because the full quantum and continuum version of the strong sector has not yet been completed. But conceptually it is one of the programme's most interesting outcomes: instead of explaining why nature tuned a knob almost exactly to zero, the proposed deeper structure suggests there may simply be no physical knob there to tune.

17. Where VERSF deliberately refuses to calculate

The neutrino sector, the third and ghostliest column of the family table, provides an important counterexample to the temptation to publish attractive numbers too soon.

Neutrinos, recall, are the extraordinarily light particles that stream through everything. They have a celebrated quirk: as they travel, they can change from one of their three types into another.

The current VERSF architecture knows how a future neutrino calculation should be performed. What it does not yet have are the fundamental physical matrices from which the neutrino masses and transformations should be calculated.

The programme went further and demonstrated something uncomfortable: if those missing matrices are simply supplied by hand, an enormous range of desired answers can be reconstructed. In other words, the framework at its current stage could be made to "predict" almost any neutrino result one wanted.

That means a close neutrino fit published today would have essentially no evidential value, and could damage the credibility of everything else.

So VERSF currently refuses to publish a numerical neutrino prediction as though it had been derived.

This is not a numerical success, but it is methodological progress: the theory has identified exactly why a premature answer would prove nothing. That refusal is as important to the standing of the programme as any of its close numbers.

18. How VERSF decides whether something has actually been derived

A recurring difficulty in foundational physics is the large gap between two statements that sound similar:

"the theory can contain this answer"

and

"the theory had no choice but to produce this answer."

Only the second is a derivation. To keep itself honest about the difference, VERSF grades every result by how far it has travelled from mathematical possibility to physical observable. In plain English:

Stage 0. We know what calculation is needed. The target is defined, but the theory has not yet produced it.

Stage 1. The underlying structure has appeared. There is a genuine precursor, but it is not yet the final physical quantity.

Stage 2. The theory returns a definite number under stated assumptions. The measured answer was not used to choose it, but some part of the physical interpretation, branch or normalisation remains unresolved.

Stage 3. The physical microscopic theory itself returns the quantity. The correct physical carrier, branch and normalisation all come from the same underlying action.

Stage 4. The microscopic result has been translated all the way to the thing an experiment measures. Only at this point is a direct experimental comparison fully meaningful.

Stage 5. All the major sectors emerge together from one compatible physical construction, rather than each from its own partial story.

VERSF also uses the intermediate label **3S**, "structural Stage 3", for cases where the physical object and mechanism are substantially identified but one crucial source-selection or normalisation step remains. It is deliberately not counted as completed Stage 3.

The stages are not decoration. They are what stops a close number at Stage 2 from being advertised as a completed derivation.

19. The rule that matters most: calculate, freeze, then compare

Behind the stages sits one simple methodological rule: a result should be built and frozen *before* the corresponding measured answer is used to judge it.

That does not mean experiment can never be used. Calibration is a legitimate scientific tool, and quantities carrying units ultimately have to be connected to metres, seconds and kilograms somehow. The point is that calibrated inputs must be labelled as inputs, and the quantities they directly determine must never be counted a second time as predictions.

In practice the programme holds itself to rules like these:

- Never choose a branch, orientation or coefficient because it gives the closest measured value.
- Never count a quantity as a prediction if it was used to calibrate the calculation.
- Keep failed outputs in the ledger; do not replace them with a better-looking alternative.
- Never compare two numbers that refer to different physical scales or conventions.
- Do not promote an attractive precursor to a physical constant before the carrier and bridge have actually been derived.
- When a later calculation supersedes an earlier one, it must be demonstrably the *same physical object*, not merely something numerically similar.

These rules do not prove that VERSF is correct. They make it much harder for the programme to fool itself. You have already seen every one of them in action in the preceding sections.

20. Where the programme actually stands today

The picture is uneven. Scientifically, that is a good sign: it means the programme is not simply manufacturing attractive numbers everywhere it looks.

Area	Where it stands	Stage	Main remaining problem
Higgs structural ratio	Exact 32/63, about 0.05% from the value inferred from measurement	2	Return it from the fully physically normalised theory
Higgs mechanism	Carrier, stable background, restoring force and fact-to-fact propagation identified	3S	The primitive theory must decide between the two spatial propagation mechanisms
Electromagnetism	7/960 candidate, about 0.08% from the measured low-energy value	2	Prove the fourteen-direction interface physically; connect and carry to the laboratory scale
Charged-lepton masses	Muon-to-electron ratio within about 0.17%; tau step within about 2.6% and more conditional	2	Derive the underlying localisation law and the exact tau mechanism
Quark mixing: baseline	A sealed calculation exists and fails badly on the held-out quantities	2	The primitive source must close before any correction can replace it
Quark mixing: correction	Unique lowest-order geometry, calculated threefold symmetry, stable full-carrier attachment	strong 3S	Same-object proof, physical readout, orientation selection, survival through commitment
Strong-force strength	Genuine precursor 1.701045 exists; the physical coupling does not	1 to 2	Complete the internal projection and the full physical reduction
Strong matter-antimatter angle	Exactly zero in the retained construction, without using the experimental bound	2 to 3	Complete the full quantum and continuum strong sector
Neutrinos	The calculation is specified; its physical inputs are absent, so no number is issued	0 to 1	Derive the physical mass and mixing inputs before any comparison
Unified boundary	The requirement that everything comes from one	1	One compatible execution across all sectors at once

Area	Where it stands	Stage	Main remaining problem
	source is defined but not yet met		

All closeness figures name the same comparator: the value measured, or inferred from measurement, by experiment.

So VERSF has not derived the Standard Model.

But neither is it merely an idea awaiting mathematics. It now contains substantial mathematical architecture, multiple microscopic constructions, several definite numerical returns, important negative results, and increasingly precise tests that the remaining theory must pass.

21. What has changed most as the programme has matured

The clearest measure of progress is not the number of pages written or equations produced. It is that the unresolved questions have become narrower.

At the beginning, many of the questions were broad. Why three families? Why does history matter? How might a particle persist? Could the Higgs emerge? Could the strength of electromagnetism be structural? Could the quark-mixing pattern arise naturally?

As the work has progressed, those broad questions have been steadily replaced by narrower ones.

The Higgs problem is no longer "find some way for a Higgs disturbance to propagate". It is now a specific yes-or-no test of which of two microscopic transport mechanisms the primitive theory actually contains.

The quark-mixing problem is no longer "find a correction that improves the data". The correction's lowest-order geometry is already derived, its threefold structure is calculated, and its attachment to the full physical carrier is executed, with the decisive future calculation's outputs stated in advance. The remaining questions are whether the full theory selects its strength, its physical readout, its discrete orientation, and its survival through the moment of commitment.

The history problem is no longer "could the past influence the future?" History-dependent probability changes have been demonstrated inside the working dynamics. The remaining problem is deriving the precise physical bridge that preserves and uses that information through renewal.

The strong-sector problem used to contain an invalid direct comparison. That comparison has been withdrawn, while a genuine precursor and an exact zero-angle result remain.

And the neutrino sector has progressed by *proving why* a premature numerical answer would be untrustworthy. A negative result of that kind narrows the path rather than decorating it.

This narrowing is probably the best single measure of how developed the programme has become. A vague research programme accumulates questions. A maturing one converts them into tests.

22. What would count as success?

VERSF should not be considered a successful derivation of the Standard Model merely because several of its numbers are close. Coincidence, selective memory and hidden fitting can each produce close numbers.

A much stronger test is required, and the programme has written it down. A successful derivation should eventually be capable of being run like a sealed experiment on the theory itself.

The primitive structure would be fixed first. Then, without looking at the measured answers, one compatible calculation would have to return:

- the strengths of the fundamental interactions;
- the Higgs properties;
- the relative elementary-particle masses;
- the quark transformation pattern and its matter-antimatter asymmetry;
- the neutrino mass and transformation structure;
- and the strong-interaction angle.

Wherever an absolute scale is still required to attach units, it would have to be clearly identified as either a single declared anchor or itself derived from the underlying structure, never smuggled in.

Those microscopic results would then have to be translated, through one consistent calculation, into the quantities laboratories actually measure: including how the interaction strengths change with energy on the way.

Only at the end would experiment be opened.

Some answers might agree. Some might fail. If the theory is truly fundamental, it does not get to change its mind after seeing which.

23. Why the failures matter as much as the successes

There is an understandable temptation in any speculative theory to advertise the best coincidences and quietly forget the bad ones. Over time, VERSF's work has increasingly institutionalised the opposite approach.

The Higgs result within about 0.05 per cent stays in the ledger. The electromagnetic result within about 0.08 per cent stays. The muon-to-electron result within about 0.17 per cent stays.

But so does the quark-mixing calculation that misses one important entry by more than 30 per cent, and the matter-antimatter measure by more than 40 per cent. The favourable corrected matrix is not allowed to replace that failure merely because it works far better.

Likewise, the programme does not print attractive neutrino predictions while the underlying inputs remain free. And it withdrew the earlier strong-force comparison when the two numbers were found to sit at incompatible scales and in incompatible conventions: an embarrassing discovery, recorded rather than buried.

This kind of bookkeeping does not prove that VERSF is right.

But without it, close numerical results would mean very little. A record that contains only successes is not a scientific record; it is an advertisement.

24. The present scientific position

The most accurate description of VERSF today is therefore neither

"VERSF has derived the Standard Model"

nor

"VERSF merely has a few interesting numerical coincidences."

The position lies between those extremes.

VERSF is now far beyond a purely verbal proposal. It has developed a substantial candidate explanation of how irreversible facts, retained history, persistent matter, interaction geometry and the three-family structure might fit together, supported by explicit microscopic constructions, repeated independent calculations, target-blind numerical outputs, impossibility results and a growing set of tests whose answers have been stated in advance.

It has produced several non-trivial numerical results from those structures before comparison with measurement. Some are unexpectedly close. Some are conditional. Some fail badly. Some sectors, especially neutrinos and the final physical strong-force coupling, remain substantially incomplete.

The most encouraging feature is not any single close number. It is that the programme is progressively replacing freedom with tests. Each time a candidate coefficient becomes a derived expression, a continuous ambiguity becomes a discrete choice, or a vague bridge becomes one explicit map that can pass or fail, the theory becomes easier to falsify. Several of the most mature sectors have now reached the point where the decisive next step is no longer inventing additional mathematics, but asking whether one primitive physical source actually selects the structures already discovered.

That is an important transition, and it is worth ending on.

A framework can be mathematically rich simply because it allows many possibilities. A fundamental theory has to do something much harder. It has to explain why *this* world, with its electrons and its two workaday quarks, its two unstable photocopied families, its 137, its stubbornly zero angle and its ghostly shape-shifting neutrinos, is the one that actually occurs, rather than all the other worlds its mathematics could have allowed.

That is now the central challenge facing the VERSF Standard Model programme.

Source basis for this general-reader synthesis

This paper is a non-technical synthesis of the current VERSF Standard Model closure corpus, with particular reliance on the following controlling or cumulative papers:

- The Standard-Model Numerical Confrontation and Provenance Theorem in VERSF (SMNC-3, 20 August 2026).
- Primitive Matter Selection in VERSF: Cumulative Executed Theorems on Fact Commitment, Record Retention and Renewal (PMS-1 cumulative, 17 August 2026).
- The Primitive Selection, History-to-Curvature and Standard-Model Attaching-Map Closure Theorem in VERSF (PSAC-1, 15 August 2026).
- The Full-Carrier Quark-Mixing Curvature Closure in VERSF (PCAC-1, 20 August 2026).

The wording here is intentionally less technical than the source papers. Where the source corpus grades a result as conditional, diagnostic, structurally incomplete or physically unadmitted, this paper preserves that status rather than upgrading it for readability.