Structural Reweighting in Rare B-Meson Decays: A VERSF Explanation of Correlated Anomalies

Pathway reweighting as a structural alternative to heavy-particle loop contributions

---

Foreword: VERSF as a real-world application

The Void Energy-Regulated Space Framework (VERSF) is most often presented as a foundational rethinking of how physical reality emerges — from discrete commitments rather than continuous spacetime. A natural question therefore arises: *can such a framework make contact with real, measurable phenomena?*

This paper addresses that question directly. We apply VERSF to one of the most persistent and structurally non-trivial anomaly sets in contemporary particle physics: rare flavour-changing neutral current (FCNC) transitions of the form $b \to s\ell^+\ell^-$.

These processes are uniquely sensitive probes of subleading dynamics. Because they are loop- and GIM-suppressed in the Standard Model (SM), even small deviations from the expected structure can produce observable effects. Over the past decade, a consistent empirical pattern has emerged:

- statistically significant deviations in exclusive branching fractions at low q^2,
- persistent distortions in angular observables such as P_5',
- and, following the 2022–2025 updates, a clear shift toward lepton-flavour-universal (LFU-preserving) anomalies.

This pattern presents a real challenge for conventional interpretations. Most beyond-the-Standard-Model (BSM) approaches attempt to explain the deviations by introducing new heavy particles — Z' bosons, leptoquarks, vector-like fermions — that modify the effective Wilson coefficients. While such models can reproduce subsets of the data, they typically introduce multiple independent parameters, have been forced into non-trivial model-building to accommodate the post-2022 LFU-preserving pattern, and — as we show — do not predict the specific four-way correlation structure that VERSF produces natively.

A structural reweighting, not new particles

The approach taken here is fundamentally different.

VERSF does not introduce new particles or modify the underlying field content. It proposes that physical amplitudes arise from sums over *admissible commitment pathways* — discrete sequences of events constrained by the closure structure of the theory. By *structural reweighting* we mean the following: the set of allowed transition histories is unchanged, but their relative contribution to the amplitude is modified by the closure structure of the substrate. No new histories are added, and no existing histories are removed — only the weight carried by each is altered. On this view:

anomalies are not generated by new interactions, but by a structural reweighting of which transition histories are most admissible.

Put differently: *the framework modifies how amplitudes are constructed, not what particles exist.*

Rare decays provide a particularly clean testing ground for this idea. Because their SM contributions are already suppressed, they amplify subleading effects in the weighting of admissible pathways. Small structural corrections — negligible in tree-level processes — can therefore become visible.

When formulated in this way, VERSF predicts a distinctive phenomenological signature — a four-feature conjunction:

- a shared kinematic shape function governing deviations across multiple operators (the feature developed as Claim A in the body),
- vector–axial linkage through a common gradient invariant, relating the shifts in the semileptonic vector and axial sectors (Claim B),
- null shifts of operators with vanishing Standard Model coefficients — right-handed currents, scalar and pseudoscalar operators — at leading order (Claim D),
- and lepton-flavour universality arising from hadronic-origin fields, as a direct consequence of the structure.

Claims D and LFU-universality are already consistent with post-2022 data; Claims A and B are the framework's sharpest falsifiable predictions, and are directly testable with existing LHCb data.

What this paper shows

The purpose of this paper is not to claim a complete replacement for the Standard Model or its effective field theory description. Rather, it demonstrates that VERSF provides a *structurally constrained* explanation of the observed anomaly pattern that does not rely on introducing new particle content. Specifically, we show that:

- the Standard Model amplitude is recovered exactly in the limit of vanishing structural correction,
- the leading-order correction arises from a uniquely constrained functional on admissible pathways,
- and the resulting Wilson-coefficient shifts exhibit a correlated structure consistent with the post-2022 experimental pattern.

In this sense, the $b \to s\ell^+\ell^-$ anomalies can be viewed as a first real-world probe of how VERSF modifies effective dynamics at low energies — and the first empirical test of whether the framework's discrete-commitment substrate produces observable signatures in the regime accessible to current experiments.

The claim, in one sentence

At its core, the claim of this paper is simple:

the observed anomalies in rare B-meson decays are not best understood as the effect of new particles, but as the result of a structured reweighting of which transition histories contribute to physical amplitudes.

The remainder of the paper shows how this idea can be made precise, and how it leads to concrete, testable predictions. In practical terms: instead of adding a new particle to the loop, the theory changes how much weight different loop-level histories carry.

This shifts the interpretation of flavour anomalies from a search for new particles to a test of how amplitudes are constructed at a more fundamental level.

Abstract

We propose that the persistent anomalies in rare $b \to s\ell^+\ell^-$ transitions are not produced by new particles in the loop, but by a structural reweighting of which loop-level transition histories are most admissible on the underlying substrate of the Void Energy-Regulated Space Framework (VERSF). Our central result is that this reweighting produces a distinctive four-feature conjunction — shared q^2-shape across operators, vector–axial linkage through a common gradient invariant, null shifts of all operators with vanishing SM coefficients, and lepton-flavour universality arising from hadronic-origin fields — and that no generic new-physics construction is known to reproduce all four features simultaneously without additional model-building assumptions. Three of the four features are already consistent with post-2022 data; the fourth is directly testable against existing LHCb data.

Formulating rare-decay amplitudes as sums over admissible commitment pathways on the K=7 closure architecture, we prove three theorems. **Theorem 1** recovers the SM effective Hamiltonian in the limit of vanishing structural correction, conditional on the substrate reproducing SM effective theory at the operator level — a scope we treat as inherited rather than derived. **Theorem 2** shows that the leading-order structural correction functional ΔS_V is uniquely determined by covariance, locality, and closure to consist of exactly three dimensionless invariants in the κ-field φ and commitment density ρ_c. **Theorem 3** derives, from pathway-class-averaged linear response, the Wilson-coefficient shift

$\delta C_i(q^2) = -C_i^{SM} \cdot \lambda_i \cdot \Xi(q^2)$,

where $\Xi(q^2)$ is a single operator-independent shape function and the λ_i are pathway-geometry ratios. This formula immediately entails the four-feature conjunction: **Claim A** (shared Ξ), **Claim B** (vector–axial linked via common $(\partial\varphi)^2$ coupling), **Claim C** (derivative-shaped q^2 profile from the gradient-dominated regime), and **Claim D** (null shift of operators with C_i^SM = 0).

Matching the 2022–2025 global-fit picture — LFU-preserving R_K, R_K*, R_φ; 4σ–5σ deviations in $\mathscr{B}(B^+ \to K^+\ell\ell)$ at low q^2; a 3.6σ deficit in $\mathscr{B}$(B_s → φμμ); and persistent P_5' tension — we identify the η_2 (gradient-dominated) regime as the VERSF-preferred solution. The empirical shift from LFU-violating δC_9^μ (pre-2022) to LFU-universal δC_9^U (post-2022), which has forced non-trivial model-building in conventional new-physics scenarios, emerges natively here: reweighting through hadronic-side fields does not couple to lepton flavour at leading order. Claim D is consistent with all current bounds on right-handed, scalar, and pseudoscalar operators at their measured precision.

The sharpest test of the pathway-ratio structure is the substrate-derived value of λ_{10}/λ_9. The simplest guess $\lambda_{10}/\lambda_9 \approx 1$ — corresponding to $\delta C_{10} \approx -\delta C_9$, the standard left-handed-leptonic NP benchmark — is excluded at several sigma by the current $\mathscr{B}$(B_s → $\mu^+\mu^-$) measurement. The framework therefore survives only if the companion PGL(3,2) calculation delivers $|\lambda_{10}/\lambda_9| \lesssim 0.2$ (axial decoupling). The q^2-shape of δC_9 inside [1.1, 6] GeV² provides a second decisive test accessible via shape-unconstrained global-fit reanalysis of existing LHCb data. The combination of these two tests, together with Upgrade-II precision on primed operators and on $B \to K^{(*)}\tau^+\tau^-$, places the framework at genuine empirical risk within the current decade.

Table of Contents

---

1. Introduction

FCNC transitions such as $b \to s\ell^+\ell^-$ occur at one loop in the SM and carry two simultaneous suppressions (loop factor and GIM cancellation). Their effective Hamiltonian [1],

$$\mathcal{H}_{\rm eff} = -(4G_F/\sqrt{2})\, V_{tb} V_{ts}^* \, \Sigma_i\, C_i(\mu)\, \mathcal{O}_i(\mu),$$

is dominated by the dipole $\mathcal{O}_7$, the vector semileptonic current $\mathcal{O}_9$, and the axial semileptonic current $\mathcal{O}_{10}$. Anomalies are customarily parameterised as

$$C_i^{\rm eff} = C_i^{\rm SM} + \delta C_i,$$

with δC_i attributed to virtual exchange of new heavy states (Z', leptoquarks, vector-like fermions).

The alternative pursued here. We do not add particles. We modify the *weighting of admissible transition pathways* in a theory whose fundamental substrate is discrete commitment. In VERSF, physical amplitudes are not fundamentally mediated by fields propagating in continuous spacetime; they are sums over histories of commitment events on a structural substrate (K=7, PGL(3,2)-irreducible) that only emerges as an effective spacetime process in the low-energy limit. The anomaly, on this view, is not produced by a new particle in the loop, but by a structural reweighting of which loop-level transition histories are most admissible.

1.1 Why rare decays are structurally sensitive

Rare flavour-changing neutral current processes are an especially natural place to look for VERSF-type corrections *because* they are already suppressed in the Standard Model. Tree-level channels are dominated by leading continuum dynamics and are therefore comparatively insensitive to small structural reweightings of admissible pathways. By contrast, loop- and GIM-suppressed transitions amplify subleading distortions in pathway weighting. In this sense, rare decays are not an arbitrary test bed for the framework: they are precisely the observables in which fine-grained admissibility structure is most likely to survive coarse-graining.

This also explains why such effects are not seen generically across all processes: in tree-level channels the dominant amplitudes wash out small structural reweightings, whereas in loop-suppressed channels these corrections are comparatively amplified.

The paper is organised as follows. Section 2 establishes notation and states the pathway postulate. Section 3 proves SM recovery. Section 4 constructs the unique leading-order correction functional. Section 5 derives the Wilson-coefficient shifts. Section 6 confronts the post-2022 LHCb picture. Section 7 lists epistemic status line by line.

---

2. Preliminaries

2.1 The pathway object

Definition 1 (Admissible pathway). A pathway Γ for the transition $b(p_b) \to s(p_s)\, \ell^+(k_+)\, \ell^-(k_-)$ is a finite sequence of commitment events

$\Gamma = (e_0, e_1, \ldots, e_N),\ e_k \in \mathbb{K}_7,$

where $\mathbb{K}_7$ denotes the K=7 closure substrate, together with a parameterisation $\lambda \in [0, N]$ along which distinguishability accumulates. A pathway is **admissible** if (i) its boundary data match the external states, (ii) each transition $e_k \to e_{k+1}$ respects the compact cyclic closure on K=7, and (iii) the accumulated distinguishability is bounded by the closure capacity. The admissible set is denoted $\mathcal{C}(b\to s\ell\ell)$.

In VERSF, a transition amplitude cannot be fundamentally represented as a sum over continuous spacetime histories, because continuous spacetime is itself emergent and not primitive. The minimal replacement is therefore a sum over admissible commitment pathways: finite ordered sequences of distinguishability-forming events whose composition respects the K=7 closure architecture. This is not an arbitrary reformulation of the path integral, but the discrete counterpart forced by the ontology of the framework. Continuous path integration is then recovered only in the continuum limit of densely accumulated admissible commitments.

[Foundational postulate] We therefore take physical amplitudes to decompose as

$\mathcal{A}(b \to s\ell\ell) = \Sigma_{\Gamma \in \mathcal{C}}\, w(\Gamma).$ (P1)

2.2 Weighting

Theorem 0 (Weighting uniqueness, given admissibility). *The unique weighting* $w : \mathcal{C} \to \mathbb{C}$ *satisfying (a) multiplicativity on concatenated pathways,* $w(\Gamma_1 \circ \Gamma_2) = w(\Gamma_1)\, w(\Gamma_2)$*, (b) local dependence on the segment* $\{e_k, e_{k+1}\}$*, and (c) compatibility with the closure operator on K=7, is*

$$w(\Gamma) = \exp[-S\_eff(\Gamma)],$$

with S_eff additive along Γ.

This is the VERSF analogue of the standard path-integral derivation; what is novel is that additivity is enforced by compositional closure on $\mathbb{K}_7$ rather than by the assumption of a continuous Lagrangian. It does *not* by itself produce corrections beyond the SM; it fixes the form of w.

The important point is that the exponential weighting is not chosen for convenience. Once admissible histories compose sequentially, and once the contribution of a composite pathway must factor through its local segments, the weighting must be multiplicative under concatenation and therefore exponential in an additive structural action. In this sense, S_eff is not an arbitrary import from continuum field theory, but the unique bookkeeping object compatible with compositional closure on the substrate.

We decompose

$$S\_eff(\Gamma) = S\_SM(\Gamma) + \Delta S\_V(\Gamma),$$

where S_SM is the standard SM effective action along Γ (recovering conventional amplitudes in the continuum limit — see §3) and ΔS_V is the structural correction (constructed in §4).

Convention on Euclidean vs Lorentzian formalism. The weighting exp[−S_eff] as written is Euclidean; in Lorentzian quantum amplitudes one would write exp[iS_eff]. We work throughout in the Euclidean formulation, with analytic continuation to physical Minkowski kinematics performed after the pathway sum is taken and before matching onto Wilson coefficients. This is consistent with the interpretation of ΔS_V as a statistical weighting over admissible configurations, rather than a unitary time-evolution phase: in VERSF, commitment events are distinguishability-forming transitions on a discrete substrate, and their accumulated weighting is naturally partition-function-like rather than phase-like. The empirically preferred new-physics solutions in current global fits correspond to *real* Wilson-coefficient shifts, which emerge directly from Euclidean weighting — a Lorentzian formulation with exp[iS] would give imaginary (absorptive, CP-odd) shifts at leading order, in conflict with data. The Wick rotation is a technical step whose consistency with the K=7 closure structure is assumed here and is discussed further in the upstream substrate papers.

2.3 VERSF field content

The scalar closure field φ entering ΔS_V is not introduced ad hoc for the present phenomenology. It is identified with the **κ-field**, the unique scalar closure mode already isolated in the upstream VERSF spectrum as the effective low-energy carrier of closure perturbations. Its mass is fixed by the closure scale,

$m_\kappa^2 = (3/4)\, \xi^{-2}$,

and the characteristic scale governing the structural correction is correspondingly

$\Lambda_V \equiv \xi^{-1}$.

Similarly, the density $\rho_c(x)$ is identified with the local commitment density of the VERSF Single-Source field $\rho(x,t)$, with ρ_* its vacuum value. These identifications are **inherited, not fitted**: the present paper does not introduce a new effective scalar sector, but applies the already-derived closure sector of VERSF to flavour observables. In an EFT with free Λ, one has an arbitrary new scale; here Λ_V is fixed by the closure-scale axioms.

Intuitively, φ encodes how the closure structure of the substrate is locally perturbed. Its gradients therefore measure how rapidly the admissibility structure of pathways changes across the effective spacetime description, making $(\partial\varphi)^2$ the natural quantity controlling deviations in pathway weighting.

2.4 Dimensions (natural units $\hbar = c = 1$)

$[\varphi] = 0$ (dimensionless closure-field phase) $[\rho_c] = [\rho_*] = 0$ *(dimensionless, normalised to* ρ_**)* $[\Lambda_V] = [\xi^{-1}] = 1$ (mass) $[\lambda] = 0$ (dimensionless tick-count parameter) $[\partial_\mu\varphi] = 1$ (spacetime gradient of dim-0 field, in the effective continuum limit) $[S_V] = 0$

Resolution of the apparent mismatch inside $\int_\Gamma d\lambda\, (\partial\varphi)^2/\Lambda_V^2$. The path parameter λ is defined as the dimensionless tick-count along Γ; the integrand $(\partial\varphi)^2/\Lambda_V^2$ is dimensionless by direct inspection; their product $d\lambda \cdot (\partial\varphi)^2/\Lambda_V^2$ is therefore dimensionless before integration. In the effective continuum limit, each tick corresponds to a proper-time interval of order Λ_V^{-1}, and the tick-to-proper-time mapping supplies a conversion factor between the discrete-sum and continuum-integral readings of $\int_\Gamma d\lambda$. We absorb this conversion into the definition of the η_n coefficients: the η_n are effective-theory coefficients in which the tick-to-proper-time Λ_V-normalisation is already included, so that the displayed action functional is manifestly dimensionless term by term with $[d\lambda] = 0$. This convention trades a dimensionally-carrying proper-time differential $d\tau$ with $[d\tau] = -1$ for a dimensionless tick differential $d\lambda$ plus a Λ_V-rescaling of η_n; the physical content is identical and the tick formulation is more natural given the discrete-substrate origin of the theory. (See Appendix A for the tabulated bookkeeping.)

---

3. Pathway-integral formulation and SM recovery

Using (P1) and Theorem 0,

$\mathcal{A} = \Sigma_\Gamma \exp[-S_{SM}(\Gamma)] \exp[-\Delta S_V(\Gamma)]$.

Theorem 1 (SM recovery). *In the limit $\Delta S\_V \to 0$, the pathway sum $\Sigma\_\Gamma \exp[-S\_SM(\Gamma)]$ reduces to the SM matrix element*

$\langle s\ell\ell|\mathscr{H}\_eff\hat{}SM|b\rangle = -(4G\_F/\sqrt{2})\ V\_tb\ V\_ts^*\ \Sigma_i\ C_i\hat{}SM\ \langle s\ell\ell|\mathcal{O}_i|b\rangle$.

Sketch of proof. Partition $\mathcal{C}$ by the topological class of the pathway — each topological class corresponds to a distinct operator insertion in the short-distance expansion at the scale μ = m_b. Closure on $\mathbb{K}_7$ in the continuum limit (accumulated distinguishability → ∞, tick parameter → proper time) is assumed to reproduce the Feynman-diagram sum of the SM effective theory. The operator basis $\{\mathcal{O}_7, \mathcal{O}_9, \mathcal{O}_{10}, \ldots\}$ is then identified with the set of pathway topologies surviving GIM cancellation at dimension six. Thus the standard effective Hamiltonian is not replaced but *recovered* as the low-energy, continuum coarse-graining of the admissible pathway sum. The VERSF correction enters only through non-trivial reweighting of those same admissible classes. ■

What this proof does not do. Theorem 1 as stated assumes, rather than derives, that the pathway-to-diagram map reproduces:

- the CKM structure V_tb V_ts* weighting the effective Hamiltonian;
- the GIM cancellation pattern specific to up-type-quark loops in the SM;
- the operator renormalisation-group running from μ = M_W down to μ = m_b.

Each of these is a non-trivial feature of SM loop dynamics and none of them is derived from the K=7 substrate in the present paper. A full substrate-level derivation of CKM structure, GIM, and RG running from pathway enumeration is an open problem in the wider VERSF programme, to be addressed in dedicated papers. In §7 we label Theorem 1 accordingly as *assumed-at-substrate-level*, not *derived-from-substrate*.

Scope. The present paper should be read as a **conditional extension**: given that the VERSF substrate reproduces SM effective theory at the operator level, we analyse how structural reweighting modifies it. This scope is deliberate — the structural-reweighting content is logically separable from the substrate-to-QFT reduction, and treating it independently is what allows the correction functional to be derived cleanly.

Remark. Theorem 1 is the *minimal* check: any framework claiming to modify the SM by reweighting must reduce to the SM when the modification vanishes. It does.

---

4. The structural correction functional

We construct ΔS_V as a scalar functional on admissible pathways, constrained by:

(S1) 4-diffeomorphism covariance (the effective spacetime symmetry that emerges from K=7 closure), (S2) locality along Γ (integrand depends on fields at the current event only, or their derivatives at that event), (S3) dimensional consistency with $[S_V] = 0$, (S4) compatibility with closure: ΔS_V vanishes on trivial (identity) pathways.

Theorem 2 (Basis of lowest-dimension structural corrections). *The complete basis of scalar functionals satisfying (S1)–(S4), containing at most two derivatives of φ and linear in $\delta\rho_c$, consists of exactly three terms:*

$$\Delta S_V(\Gamma) = \eta_1 \int_\Gamma d\lambda\, (\delta\rho_c / \rho^*) + (\eta_2 / \Lambda_V^2) \int_\Gamma d\lambda\, \partial_\mu \varphi\, \partial^\mu \varphi + (\eta_3 / \Lambda_V^2) \int_\Gamma d\lambda\, \Box\varphi.$$

Proof sketch. Enumerate scalar monomials built from {φ, *∂μφ, □φ, δρ_c/ρ} with total mass dimension zero after contraction and integration. At dimension zero in derivatives: only δρ_c/ρ_* (φ itself is pure phase and drops out by (S4)). At dimension two in derivatives: $(\partial\varphi)^2$ and □φ, each carrying a factor Λ_V^{-2}. All higher-order invariants carry additional powers of Λ_V^{-2} and are suppressed. ■

In particular, the absence of a free potential term V(φ) at this order is not a choice but a consequence of closure compatibility. Condition (S4) requires that ΔS_V vanish on identity pathways — pathways that make no progress through the K=7 closure cycle. A pure potential term V(φ), integrated along any path, depends only on the value of φ at each point and not on the pathway's engagement with the closure structure; in particular, it assigns the same weight to identity and non-identity pathways of equal parameter length. Such a term therefore cannot reflect the substrate-level admissibility it is supposed to encode, and is excluded by (S4). The leading structural correction must therefore begin with density *departure* ($\delta\rho_c$, which vanishes when no commitment progress occurs) and derivative *sensitivity* ($\partial_\mu\varphi$, which vanishes on stationary paths) — exactly as in the basis above.

Physical interpretation:

- η_1 — sensitivity to local departures of commitment density from vacuum.
- η_2 — sensitivity to gradients of the closure field (kinetic term).
- η_3 — sensitivity to the Laplacian of the closure field (relaxation / curvature).

The key point is that the allowed correction is not freely chosen: it is the lowest-order functional compatible with closure, locality, and vanishing on identity pathways. Any other correction at this order either violates one of these conditions or reduces to a combination of the three terms above. The framework's predictive power rests precisely on this inevitability.

[Status] Theorem 2 is **derived** under the stated symmetry and derivative-count restrictions. Extending the basis to quadratic in $\delta\rho_c$ or to four-derivative φ terms introduces further η_i at higher order in $(E/\Lambda_V)^2$; at b-quark scales $E/\Lambda_V \ll 1$ and these are negligible provided $\xi^{-1} \gtrsim$ 10 GeV. This is the one **assumed** input of the phenomenology: the closure scale sits above the b mass.

Quantitative validity of the leading-order truncation. The natural expansion parameter is $(m_b / \Lambda_V)^2$. With $m_b \approx 4.8$ GeV and $\Lambda_V = \xi^{-1}$, a closure scale of $\Lambda_V \approx 10$ GeV gives $(m_b / \Lambda_V)^2 \approx 0.23$. This is *not* a small number: it is of the same order as the claimed Wilson-coefficient shift $\delta C_9 / C_9^{SM} \approx 0.23$. The leading-order truncation therefore cannot be justified purely by power counting at b-meson scales — next-to-leading-order corrections are expected at the 20–25% level of the leading result, which is the same size as the effect being predicted. The leading-order expressions below should therefore be read as a *first approximation* to a ~25% NLO uncertainty, not as a parametrically controlled prediction. A firm quantitative prediction at the 10% level requires either (i) pushing the framework to NLO, or (ii) an independent determination of Λ_V placing it substantially above 10 GeV. Both are open problems.

Structural vs magnitude predictions. This ~25% truncation uncertainty affects the *magnitude* of δC_9 predicted by the framework, but not the *structural* content — the shared shape function $\Xi(q^2)$, the pathway-geometry ratios λ_i, and the null-shift of vanishing-SM-coefficient operators (Claims A, B, D). These structural features are preserved under NLO corrections, which modify the overall normalisation of the correction but not its operator-level topology. Quantitative fits to the magnitude of δC_9 from this framework should therefore be treated as order-of-magnitude until NLO corrections are computed; the structural predictions, by contrast, are robust at leading order.

---

5. Wilson-coefficient projection

This is the central calculation, and the step that was hand-waved in the first draft.

5.1 Operator-graded decomposition of $\mathcal{C}$

Partition admissible pathways by topological class:

$\mathcal{C} = \sqcup_i \mathcal{C}_i$, with $\mathcal{C}_i$ = {pathways contributing to $\mathcal{O}_i$ at scale $\mu = m_b$}.

By Theorem 1,

$\mathcal{A} = \Sigma_i C_i^{SM} \cdot \mathcal{U}_i$, with $\mathcal{U}_i \equiv \Sigma_{\Gamma \in \mathcal{C}_i} \exp[-S_{SM}(\Gamma)] / \langle s\ell\ell|\mathcal{O}_i|b\rangle \cdot \langle s\ell\ell|\mathcal{O}_i|b\rangle$,

so that $\mathcal{U}_i = \langle s\ell\ell|\mathcal{O}_i|b\rangle$ in the $\Delta S_V \to 0$ limit.

5.2 Linear-response correction

To leading order in ΔS_V,

$\exp[-S_{SM} - \Delta S_V] \approx \exp[-S_{SM}] \cdot (1 - \Delta S_V)$.

Therefore,

$\mathcal{A}_{VERSF} = \mathcal{A}_{SM} - \Sigma_i C_i^{SM} \cdot \langle \Delta S_V \rangle_i \cdot \langle s\ell\ell|\mathcal{O}_i|b\rangle$, where

$$\langle \Delta S_V \rangle_i \equiv [\, \Sigma_{\Gamma \in \mathcal{C}_i} \exp[-S_{SM}(\Gamma)]\, \Delta S_V(\Gamma)\,] \,/\, [\, \Sigma_{\Gamma \in \mathcal{C}_i} \exp[-S_{SM}(\Gamma)]\,].$$

This is a pathway-class-weighted expectation value of ΔS_V — a well-defined, calculable object once $\mathcal{C}_i$ is specified.

Validity of linear response. The expansion above is controlled by $\langle \Delta S_V \rangle_i \ll 1$. Given that the claimed phenomenological effect is $\delta C_9/C_9^{SM} \approx 0.23$, we have $\langle \Delta S_V \rangle_9 \approx 0.23$ — on the edge of the linear regime but not deep inside it. Quadratic corrections ($\frac{1}{2}\langle \Delta S_V^2 \rangle_i$) are expected at ~3% of the leading SM amplitude, or ~12% of the leading δC_9 shift (i.e. a ~12% correction to the magnitude of δC_9 itself, not to C_9^{SM}). These are the same order as the NLO truncation corrections noted in §4 and should be kept in mind when comparing to data at the 10% level.

5.3 Derivation of δC_i

Comparing with the standard parameterisation $\delta C_i \langle \mathcal{O}_i \rangle$,

$$\boldsymbol{\delta C_i = -C_i^{SM} \cdot \langle \Delta S_V \rangle_i.} \quad \text{(Theorem 3)}$$

This is **derived**, not postulated.

The significance of Theorem 3 is that the Wilson-coefficient shift is not inserted at the operator level. It emerges from pathway-class averaging of the structural correction. In other words, the effective short-distance anomaly is a coarse-grained reflection of substrate-level admissibility reweighting: the anomaly is not produced by a new particle in the loop, but by a structural reweighting of which loop-level transition histories are most admissible.

5.4 Shared shape function and operator-specific ratios

Because ΔS_V has three structural terms (Theorem 2), $\langle \Delta S_V \rangle_i$ decomposes as

$$\langle \Delta S_V \rangle_i = \alpha_i^{(1)} \langle \delta\rho_c/\rho_* \rangle_\Gamma + \alpha_i^{(2)} \Lambda_V^{-2} \langle (\partial\varphi)^2 \rangle_\Gamma + \alpha_i^{(3)} \Lambda_V^{-2} \langle \Box\varphi \rangle_\Gamma,$$

where $\alpha_i^{(n)}$ depend on how the pathway geometry of $\mathcal{C}_i$ couples to each structural term, and the bracketed expectation values are **operator-independent** functions of the external kinematics (q^2, m_b, hadronic form-factor convolutions).

Introduce the shape function

$$\Xi(q^2) \equiv \alpha_V \langle \delta\rho_c/\rho_* \rangle(q^2) + \beta_V \Lambda_V^{-2} \langle (\partial\varphi)^2 \rangle(q^2) + \gamma_V \Lambda_V^{-2} \langle \Box\varphi \rangle(q^2),$$

with $(\alpha_V, \beta_V, \gamma_V) = (\eta_1, \eta_2, \eta_3)$ rescaled by a common pathway normalisation. Then

$$\boldsymbol{\delta C_i(q^2) = -C_i^{SM} \cdot \lambda_i \cdot \Xi(q^2)}, \quad \lambda_i \equiv \langle \Delta S_V \rangle_i / \Xi(q^2).$$

The λ_i are **ratios** — they do not introduce new free parameters once one $\langle \Delta S\_V \rangle_i$ is known. They are determined by the operator pathway geometry, as follows.

5.5 Predicted operator ratios

Each Wilson-operator pathway class $\mathcal{C}_i$ couples differently to the three structural terms:

Operator	Pathway feature	Dominant structural coupling
$\mathcal{O}_7$ (dipole)	short-distance photon emission; local field-strength insertion	η_3 ($\Box\varphi$, curvature); weak coupling to η_2
$\mathcal{O}_9$ (V)	semileptonic vector current; vector coupling over finite q^2 range	η_2 ($(\partial\varphi)^2$)
$\mathcal{O}_{10}$ (A)	semileptonic axial current; pathway class probes gradient sector	η_2, with coefficient λ_{10} of pathway-geometry-determined magnitude relative to λ_9

Importantly, the sign of δC_i is set by the combination $-C_i$^SM · λ_i. In the SM one has, at μ = m_b, C_9^SM ≈ +4.1 and C_{10}^SM ≈ −4.3 — opposite signs. The sign relationship between δC_9 and δC_{10} is therefore controlled by the *product* (C_{10}^SM/C_9^SM) · (λ_{10}/λ_9), not by λ_{10}/λ_9 alone. This is important for the correct reading of Claim B in §5.6.

[Partially derived; substrate closure incomplete] The qualitative hierarchy of operator couplings follows already from the pathway-class structure: $\mathcal{O}_9$ and $\mathcal{O}_{10}$ both probe the gradient sector (η_2), while $\mathcal{O}_7$ is more sensitive to the curvature sector (η_3). What remains unfinished is the substrate-level derivation of the exact numerical ratios $\lambda_7 : \lambda_9 : \lambda_{10}$ from the PGL(3,2) chirality assignment on K=7. This is not a secondary detail but the **principal open calculation of the programme**. If completed, it would remove the last phenomenological freedom at leading order and convert the present framework from a structurally constrained effective correction into a first-principles substrate prediction.

Why axial suppression is structurally plausible. Although the precise value of λ_{10}/λ_9 awaits the companion calculation, a qualitative expectation can be extracted from the chirality structure already present in the Standard Model at the matching scale. In the chiral basis, $\mathcal{O}$_L ≡ ($\mathcal{O}_9$ − $\mathcal{O}_{10}$)/2 couples only to left-handed leptons and $\mathcal{O}$_R ≡ ($\mathcal{O}_9$ + $\mathcal{O}_{10}$)/2 only to right-handed leptons; the SM at μ = m_b has C_L^SM ≈ +4.2 and C_R^SM ≈ −0.1, a near-total dominance of left-handed coupling inherited from the SU(2)_L structure of the electroweak penguin loop. If the PGL(3,2) action on admissible pathways inherits this chirality structure at leading order — which is the minimal expectation compatible with Theorem 1's assumption that the substrate reproduces SM effective theory — then the axial combination λ_L − λ_R will be substantially smaller than the vector combination λ_L + λ_R, and $\lambda_{10}/\lambda_9 < 1$ follows. Whether the suppression reaches the data-required level $|\lambda_{10}/\lambda_9| \lesssim 0.2$, and whether any additional kinematic suppression enters through the pathway-averaging of $(\partial\varphi)^2$, are questions the companion PGL(3,2) calculation must settle. What we can say now is that axial decoupling is a natural structural expectation, not an unexplained coincidence.

Honest parameter count. Until the companion paper derives the λ_i ratios, the framework at leading order has:

- one functional degree of freedom — the shape $\Xi(q^2)$, constrained by structural covariance and locality;
- three dimensionless ratios — λ_7, λ_9, λ_{10}.

A generic new-physics global fit has three free functions $\delta C_7(q^2)$, $\delta C_9(q^2)$, $\delta C_{10}(q^2)$. The framework therefore reduces functional freedom by a factor of three, at the cost of a three-parameter dimensionless ratio structure whose derivation is the principal open calculation. This is a substantial reduction, but it is not (yet) a parameter-free prediction.

5.6 Structural signature

The structural content of the framework at leading order breaks into four claims of differing strengths.

(Claim A) — shared shape. The leading VERSF contribution does not generate three independent Wilson-coefficient deformations. It generates a single structural profile $\Xi(q^2)$, seen through three operator classes. At fixed λ_i ratios, a generic new-physics global fit has three free functions; VERSF has one.

(Claim B) — vector–axial linkage via λ_{10}/λ_9. The axial and vector semileptonic pathway classes both probe the gradient invariant $(\partial\varphi)^2$. Their shifts are therefore linked:

$\delta C_{10} / \delta C_9 = (C_{10}^{SM} / C_9^{SM}) \cdot (\lambda_{10} / \lambda_9) \approx (-1) \cdot (\lambda_{10} / \lambda_9)$.

The naïve guess $\lambda_{10} \approx +\lambda_9$ (equal gradient coupling — the simplest guess compatible with both classes probing the same invariant) yields the left-handed-leptonic benchmark $\delta C_{10} \approx -\delta C_9$, a well-known new-physics scenario realised in SU(2)*L-invariant couplings to left-handed leptons (S_3 or U_1 leptoquarks, certain Z' models, the SMEFT combination* $\mathcal{O}\{\ell q\}^{(1)} - \mathcal{O}_{\{\ell q\}}^{(3)}$). We emphasise that this relation in isolation is **not** VERSF-distinctive; what is distinctive is the conjunction of Claim A (shared Ξ), Claim B (linked vector–axial), and Claim D below.

Equally important, the relation $\delta C_{10} \approx -\delta C_9$ with $|\delta C_9| \sim 1$ is in tension with current data (§6.3, P2). What VERSF actually predicts is the *ratio* λ_{10} / λ_9 — a single number, computable in principle from PGL(3,2) chirality. Current global-fit preferences point toward $\lambda_{10} / \lambda_9 \approx 0$ (axial decoupling) rather than $\lambda_{10} / \lambda_9 \approx +1$ (equal coupling). The companion substrate-level calculation must land on something close to the data-preferred value, or the framework at leading order is falsified.

(Claim C) — derivative-shaped q^2 profile (motivated and structurally expected, though not yet computed from the substrate). If the η_2 term dominates ΔS_V, the correction is governed by $\langle(\partial\varphi)^2\rangle(q^2)$. A straightforward qualitative argument (§6.2) suggests this quantity is largest in the low-q^2 window. A full evaluation of $\langle(\partial\varphi)^2\rangle(q^2)$ along admissible hadronic pathways is not attempted here and is left for numerical substrate analysis; Claim C at this stage is a structurally expected pattern consistent with the observed low-q^2 localisation, not a computed shape.

(Claim D) — no leading-order shift in operators with vanishing SM coefficient. Because $\delta C_i = -C_i^{SM} \cdot \langle \Delta S\_V \rangle_i$ (Theorem 3), any operator with $C_i^{SM} = 0$ at the matching scale receives no shift at leading order. This predicts, at this order:

- no shift in right-handed currents $\mathcal{O}_9'$, $\mathcal{O}_{10}'$,
- no shift in scalar/pseudoscalar operators $\mathcal{O}\_S$, $\mathcal{O}\_P$.

An analogous null result applies to the tensor operator $\mathcal{O}\_T$ (which also has vanishing SM matching coefficient at dimension six), though the current experimental access to $\mathcal{O}\_T$ is weaker and we do not treat it as a primary test channel.

The current measurement of $\mathscr{B}(B\_s \to \mu^+\mu^-)$ agrees with SM at ~10% precision (a tight constraint on $\mathcal{O}\_S$, $\mathcal{O}\_P$), and angular analyses of $B \to K^*\ell\ell$ place tight bounds on the primed operators. Claim D is therefore **consistent with current data** — and distinguishes VERSF from several leptoquark scenarios that naturally predict non-zero contributions to these operators.

The sharpest VERSF-distinctive content is the **conjunction** of these claims: one shape function, controlling the three SM-active operators with fixed pathway ratios, while leaving SM-inactive operators unshifted. No generic new-physics construction produces all four features simultaneously without additional model-building assumptions.

---

6. Confrontation with experimental observations

The key idea carried into the data comparison is that the anomaly arises from a reweighting of admissible transition histories, not from the addition of new interaction terms.

The $b \to s\ell\ell$ landscape has evolved substantially since 2022, and the current pattern is — we will argue — structurally distinctive in a way that maps cleanly onto the VERSF prediction. We therefore work through the measurements individually before stating the confrontation.

6.1 Current experimental picture (2022–2025)

Lepton-universality ratios. The LHCb 2022 update of R_K and R_K* returned values consistent with the SM expectation of unity [2], and the 2024 LHCb measurement of R_φ in $B\_s \to \varphi\ell^+\ell^-$ [3] likewise found agreement with lepton-universality across three q^2 bins. The 2025 LHCb measurement at large dilepton mass [4] reports $R\_K = 1.08\,^{+0.11}_{-0.09}$ (stat) ± 0.04 (syst) above the $\psi(2S)$ mass, again consistent with unity. Taken together, the LFU-*violating* scenarios that dominated the 2015–2021 literature are no longer favoured.

Exclusive branching fractions. The tension has migrated into the exclusive branching fractions themselves — in *both* electronic and muonic modes. $\mathscr{B}(B^+ \to K^+ \mu^+\mu^-)$ and $\mathscr{B}(B^+ \to K^+ e^+e^-)$

both deviate from SM at the 4σ–5σ level in the low-q^2 region ($q^2 \in [1.1, 6]$ GeV^2) [5]. $\mathscr{B}(B_s \to \varphi\mu^+\mu^-)$ lies 3.6σ below SM in the same q^2 window [6].

Angular observables. The angular observable P_5' in $B \to K^*\mu^+\mu^-$ shows a 3.3σ deviation from SM [7,8]. CMS 2024 angular-analysis results agree with LHCb at the 1σ level, providing independent confirmation [9]. The LHCb 2025 angular update is in agreement with previous LHCb results and the recent CMS data, but exhibits even larger tensions due to smaller experimental uncertainties [5].

Global fits. Multiple independent global-fit groups now converge on a common interpretation [5,10]. The deviations can be consistently accommodated by new-physics contributions to the effective Wilson coefficient C_9, corresponding to a modification of the semileptonic vector contribution, with best-fit value $\delta C_9 \approx -1$ at significance $\geq 4\sigma$. Critically, the preferred scenario has **shifted from LFU-violating** δC_9^μ **to LFU-universal** δC_9^U: after the 2022 R_K updates, a lepton-flavour-universal contribution to C_9 gives a better fit than a purely muonic one [11].

The NP-vs-hadronic question. The persisting ambiguity in the literature is whether these deviations reflect genuine physics beyond the SM or underestimated non-factorisable hadronic power corrections in the SM prediction itself. Eliminating the NP preference in global fits requires power corrections of a size that most authors judge unrealistically large, but this is a theory-systematics question that will only be resolved by LHCb Upgrade II.

6.2 The structural signature, matched to data

The surviving empirical pattern is, compactly:

a q^2-localised, LFU-preserving, C_9-dominant deviation appearing simultaneously across several exclusive channels, with correlated angular-observable distortions, consistent with a single underlying kinematic shape.

Each feature maps onto the VERSF framework:

(i) LFU-preservation ↔ the pathway-independence of ΔS_V from lepton flavour. ΔS_V depends on the closure field φ and the commitment density ρ_c — objects defined on the hadronic/substrate side of the transition, with no coupling to lepton flavour. At leading order the framework predicts δC_i to be lepton-flavour-universal, with deviations only at $\mathcal{O}(m_\ell^2/m_b^2)$. This matches the R_K, R_{K^*}, R_φ results directly and *without additional assumptions*. A generic new-physics model has to explain why it is LFU-preserving; VERSF has to explain why it would not be.

(ii) C_9-dominance ↔ η_2-dominance in ΔS_V. The pathway class $\mathcal{C}_9$ couples to the gradient invariant $(\partial\varphi)^2$ through the semileptonic vector current. The pathway class $\mathcal{C}_7$ couples primarily through $\Box\varphi$ (curvature of the closure field, relevant to short-distance dipole emission). The overlap coefficients $\alpha_i^{(n)}$ therefore predict the dominant shift to sit in δC_9 when η_2 dominates — which, given the structural scaling $\langle(\partial\varphi)^2\rangle / \Lambda_V^2$ versus $\langle\delta\rho_c/\rho_*\rangle$, is the expected regime for processes probing the closure field at $E \sim m_b \ll \Lambda_V$ but non-negligible on that scale. This is not a post-hoc fit; it is the natural hierarchy when the gradient term is not suppressed by additional factors of E/Λ_V.

(iii) q^2-localisation ↔ kinematic support of the gradient sector. A derivative-dominated ΔS_V produces a non-flat q^2 profile because the pathway-averaged quantity $\langle(\partial\varphi)^2\rangle(q^2)$ is sensitive to how rapidly the effective closure field must vary across the admissible hadronic transition region. In the b-meson channels of interest, this variation is maximised in the low-q^2 window because this is where the transition probes the largest relative change in the effective closure field over the accessible kinematic scale. The resulting structural correction is therefore naturally concentrated in the [1.1, 6] GeV^2 region rather than approximately flat across q^2, as would be more typical of heavy-particle contact-like mediation away from poles.

(iv) Correlated channels ↔ single Ξ controlling many observables. The same $\Xi(q^2)$ drives the shift in $B^+ \to K^+\mu\mu$, $B\_s \to \varphi\mu\mu$, $B \to K^*\mu\mu$, and (at higher q^2) $\Lambda\_b \to \Lambda\mu\mu$. The consistency of the deviation pattern across these channels, with similar q^2-shape and consistent sign, is exactly the structural correlation predicted by (Claim A).

6.3 Specific predictions sharpened by current data

(P1) LFU-universal δC_9 in the electron–muon sector. VERSF predicts δC_9^e ≈ δC_9^μ to $\mathcal{O}(m\_\mu^2/m\_b^2) \approx 5 \times 10^{-4}$ — essentially LFU-exact. This is substantially stronger than "consistent with LFU-preservation": it predicts a specific precision floor at which any residual deviation must first appear. Current R_K precision is ≈ 4% in Run-2 data; projected Upgrade-II sensitivity is ≈ 1%; the framework's predicted deviation sits roughly two orders of magnitude below either. A flavour-universal Z′ or leptoquark gives LFU-preservation with natural corrections of similar size, but typically accompanied by additional flavour-structure parameters that can be tuned. VERSF gives LFU-preservation with no such freedom at leading order: **if R_K deviates from unity by more than ~0.1% at Upgrade-II precision, the framework is falsified.** Status: passed at current precision and effectively unfalsifiable on LFU grounds for the foreseeable experimental future.

(P2) Vector–axial ratio $\delta C_{10}/\delta C_9$ set by λ_{10}/λ_9 — simplest guess excluded. Claim B predicts that $\delta C_{10}/\delta C_9 = (C_{10}\text{^SM}/C_9\text{^SM})(\lambda_{10}/\lambda_9) \approx -(\lambda_{10}/\lambda_9)$. Current post-2022 global fits [5, 10, 11] favour $\delta C_9 \approx -1$ with δC_{10} consistent with zero at the 1σ level, corresponding to the empirical ratio $\lambda_{10}/\lambda_9 \approx 0$. The simplest guess $\lambda_{10}/\lambda_9 \approx +1$ — giving the left-handed-leptonic benchmark $\delta C_{10} \approx -\delta C_9 \approx +1$ — would shift $\mathscr{B}(B\_s \to \mu^+\mu^-)$ downward by ~40%, which is *excluded at several sigma* by the current measurement (which agrees with SM at ~10% precision). Translating the $B\_s \to \mu\mu$ constraint into a bound on δC_{10} (roughly $|\delta C_{10}| \lesssim 0.15$–0.20 given the current measurement precision) and using $|\delta C_9| \approx 1$ yields the operational target

$$|\lambda_{10} / \lambda_9| \lesssim 0.2$$

which the companion PGL(3,2) calculation must satisfy. **Status: the framework is consistent with data only if the substrate calculation produces $|\lambda_{10}/\lambda_9| \lesssim 0.2$.** If the substrate calculation instead forces $\lambda_{10}/\lambda_9 \approx +1$, the framework at leading order is falsified.

(P3) q^2-shape of δC_9 — motivated and structurally expected. VERSF *indicates* $\delta C_9(q^2)$ tracks $\langle(\partial\varphi)^2\rangle(q^2)$, with qualitative arguments (§6.2 iii) pointing to low-q^2 concentration. Current global fits typically *assume* q^2-independent δC_9 — they cannot distinguish this from a smoothly q^2-varying one. A bin-by-bin analysis of δC_9 across [1.1, 6] GeV^2 would directly test whether the data prefers a shape or a constant; the 2025 LHCb data has the statistical power for this. A null result (flat δC_9) would disfavour the η_2-dominated regime.

(P4) $b \to s\tau\tau$ enhancement. A LFU-universal δC_9^U at the currently preferred magnitude (≈ -1) predicts a 20–30% enhancement of $\mathscr{B}(B \to K^{(*)}\tau^+\tau^-)$ relative to SM, concentrated at high q^2 (consistent with standard δC_9^U projections; see e.g. [10]). We emphasise that this prediction is inherited from the global-fit preference for LFU-universal δC_9 and is *not* VERSF-distinctive. The VERSF-distinctive question concerns τ-vs-μ pathway averages: at q^2 approaching $m_\tau^2 \approx 3.2$ GeV^2, the kinematic region probed by the τ channel differs from the muonic one, and $\langle(\partial\varphi)^2\rangle(q^2)$ may differ appreciably between the two — introducing a VERSF-predicted deviation from exact lepton-universality in τ-mode observables. Whether this is at the 1% or 30% level cannot be stated without evaluating $\langle(\partial\varphi)^2\rangle(q^2)$, which is an open calculation; the qualitative prediction is that τ modes should show some pathway-kinematic departure from the eμ pattern, quantifiable once the shape function is computed.

(P5) Structural correlation across channels. The ratio of deviations in $B^+ \to K^+\mu\mu$ vs $B_s \to \varphi\mu\mu$ vs $B \to K^*\mu\mu$ follows from the single $\Xi(q^2)$ weighted by channel-specific hadronic form-factors. Currently these deviations are fit independently; VERSF predicts they lie on a one-parameter locus (parameterised by the overall normalisation of Ξ). This is testable in existing global-fit infrastructure.

6.3.5 Constraints the framework already passes, with null-shift predictions

Three classes of observables tightly bound any new-physics contribution to rare $b \to s$ transitions and deserve explicit confrontation: radiative modes (constraining δC_7), the $B_s \to \mu\mu$ branching ratio (constraining δC_{10} and the scalar/pseudoscalar sector), and angular analyses of $B \to K^*\ell\ell$ (constraining primed operators). The first two constrain the pathway-ratio structure; the third is where Claim D — the framework's sharpest null prediction — most directly meets data.

$\mathscr{B}(B \to X_s\gamma)$ and $\mathscr{B}(B \to K^*\gamma)$. These radiative modes are sensitive to δC_7 and currently constrain $|\delta C_7| \lesssim 0.05$ (order-of-magnitude; the precise HFLAV-averaged bound at 1σ lies in the range 0.03–0.05 depending on fit and scale). Under Theorem 3, $\delta C_7 = -C_7^{SM} \cdot \lambda_7 \cdot \Xi$ with $C_7^{SM} \approx -0.3$, while $\delta C_9 = -C_9^{SM} \cdot \lambda_9 \cdot \Xi$ with $C_9^{SM} \approx +4.1$. The ratio is therefore

$$\delta C_7/\delta C_9 = (C_7^{SM}/C_9^{SM}) \cdot (\lambda_7/\lambda_9) \approx -0.073 \cdot (\lambda_7/\lambda_9).$$

With $|\delta C_9| \approx 1$, requiring $|\delta C_7| \lesssim 0.05$ gives $|\lambda_7/\lambda_9| \lesssim 0.68$. This is a *non-trivial constraint on the pathway-ratio structure* that must be satisfied by the companion PGL(3,2) calculation. It is consistent with the expectation (§5.5) that $\mathcal{O}_7$ couples primarily to η_3 (curvature) rather than η_2 (gradient), since η_3-mediated contributions are suppressed relative to η_2-mediated ones when the η_2-dominant regime holds. The framework's internal consistency therefore *requires* that λ_7 be substantially smaller than λ_9, as expected — but this is a testable internal check, not an automatic success.

$\mathscr{B}$(B_s → $\mu^+\mu^-$). Dominated by $\mathcal{O}_{10}$ and sensitive to $\mathcal{O}$_S, $\mathcal{O}$_P, the current measurement agrees with SM at ~10% precision. Claim D (null shift of $\mathcal{O}$_S, $\mathcal{O}$_P at leading order) is automatically consistent with this. Claim B in its simplest form ($\delta C_{10} \approx -\delta C_9 \approx +1$) would shift $\mathscr{B}$(B_s → μμ) downward by ~40% and is *excluded at several sigma*; the framework survives this constraint only if $|\lambda_{10}/\lambda_9| \lesssim 0.2$ (see P2 for the derivation of this bound). The substrate-level PGL(3,2) calculation is therefore not an unconstrained open question: it must deliver $|\lambda_{10}/\lambda_9|$ below this operational target, or the leading-order framework is falsified.

Null-shift predictions confronted with current bounds. Claim D states that every Wilson operator with vanishing SM coefficient at the matching scale receives zero VERSF shift at leading order. This is a sharp, multi-operator prediction: the framework produces no contribution to the right-handed currents or the scalar/pseudoscalar operators, regardless of the detailed pathway geometry. We compile the current experimental bounds against this prediction in Table 1.

Table 1. Null-shift predictions from Claim D vs current experimental bounds on SM-inactive Wilson coefficients. All bounds are 1σ and reflect typical post-2022 fit ranges; specific numerical values vary between global-fit groups (e.g. [5, 10, 11]). Upgrade-II projections are expected reach based on luminosity scaling; they should be regarded as order-of-magnitude.

Operator	SM coefficient at μ = m_b	VERSF prediction	Current bound (1σ, typical)	Upgrade-II expected reach	Consistency
$\mathcal{O}_9'$ (right-handed V)	≈ 0	$\delta C_9' = 0$	$\|\delta C_9'\| \lesssim 0.4$	~0.1	✓
$\mathcal{O}_{10}'$ (right-handed A)	≈ 0	$\delta C_{10}' = 0$	$\|\delta C_{10}'\| \lesssim 0.4$	~0.1	✓
$\mathcal{O}$_S (scalar)	≈ 0	δC_S = 0	$\|\delta C_S\| \lesssim 0.02$ (from $\mathscr{B}$(B_s → μμ))	factor ~2 improvement expected	✓
$\mathcal{O}$_P (pseudoscalar)	≈ 0	δC_P = 0	$\|\delta C_P\| \lesssim 0.02$ (from $\mathscr{B}$(B_s → μμ))	factor ~2 improvement expected	✓

The framework passes all four constraints at leading order. (We do not tabulate the tensor operator $\mathcal{O}$_T: its experimental access is weaker and less commonly quoted, and we do not treat it as a primary test channel.) Several features of the table deserve emphasis.

First, the $\mathcal{O}$_S and $\mathcal{O}$_P bounds from $\mathscr{B}$(B_s → $\mu^+\mu^-$) are the tightest, and they are tight precisely because these operators would produce helicity-unsuppressed contributions to a helicity-suppressed SM amplitude — the measurement is therefore exquisitely sensitive. Claim D's prediction of exact null shift at leading order is automatically consistent with this precision, and will remain consistent through Upgrade-II sensitivity improvements unless next-to-leading-order

VERSF corrections exceed the ~5% level. The latter would correspond to $(m\_b/\Lambda\_V)^4 \approx 0.053$ contributions from NLO terms not present in the Theorem 2 basis; this is possible in principle but would itself constitute a derivable prediction rather than a fit parameter.

Second, the right-handed-current bounds are significantly weaker than the scalar ones, and the framework's prediction of $\delta C_9' = \delta C_{10}' = 0$ is consistent with both the current fits and the projected Upgrade-II reach. If a future angular analysis finds $\delta C_9'$ or $\delta C_{10}'$ at the $|\delta C_i'| \sim 0.1$–0.3 level, Claim D at leading order is falsified; absorbing such a shift into higher-order VERSF corrections would require unnaturally large NLO contributions, and is disfavoured as a rescue strategy.

Third, Table 1 makes the contrast with leptoquark-based NP explanations concrete. The S_3 leptoquark naturally generates $\delta C_9'$ and $\delta C_{10}'$ at the same magnitude as δC_9 and δC_{10}; the R_2 leptoquark generates non-zero δC_S and δC_P; the U_1 vector leptoquark can avoid primed operators but generates specific correlated shifts (for a comprehensive review of leptoquark phenomenology see [17]). The VERSF prediction of simultaneous null shift across the four tabulated SM-inactive operators is therefore structurally different from the major leptoquark scenarios, and the conjunction of this null-shift pattern with the gradient-dominated δC_9 of (P2) is the feature no single NP model reproduces.

The four passed constraints do not, of course, *demonstrate* VERSF — they are passed automatically by any theory that happens to generate only δC_9-type shifts, including pure hadronic power corrections. What they demonstrate is that the framework's structural prediction (Claim D) is consistent with the most sensitive current measurements, and that the framework remains viable at Upgrade-II precision. Falsification of Claim D by future Upgrade-II data would be a clean failure mode, distinguishable from NLO corrections, and is the leading-order framework's most directly testable null prediction.

Collectively, these constraints are non-trivial: they say that a viable VERSF realisation must have $|\lambda_7/\lambda_9| \lesssim 0.7$ (from $b \to s\gamma$), $|\lambda_{10}/\lambda_9| \lesssim 0.2$ (from $B\_s \to \mu\mu$), and zero leading-order shift in primed and scalar/pseudoscalar operators (automatic from Theorem 3). The first two are numerical targets the companion PGL(3,2) calculation must deliver; the third is the leading-order prediction the Upgrade-II programme will most directly test.

6.4 Position relative to the hadronic-vs-NP debate

The framework is *not* in competition with the hadronic-corrections interpretation — it is structurally orthogonal to it. If the persistent deviations turn out to be dominantly hadronic, VERSF predicts a residual structural correction at reduced significance, testable with Upgrade-II precision. If they are dominantly new physics, VERSF predicts the new physics has the specific structural form of (P1)–(P5) and §6.3.5. The hadronic-vs-NP question is about the *size* of the VERSF contribution; the framework's distinctive predictions concern its *structure*.

Crucially, many conventional new-physics models (Z′, leptoquarks, loop-induced operators) now face tension with the LFU-preservation data: they were originally motivated by the LFU-violating picture and require non-trivial model-building to accommodate the current LFU-preserving pattern. VERSF accommodates it natively — LFU-preservation is a consequence of the structure, not an imposed feature.

6.5 Benchmark regime selection

We adopt as a phenomenological working regime:

$\eta_2 \gg \eta_1, \eta_3$

This is selected empirically, not derived: it is the hierarchy that best accommodates (i) the observed low-q^2 concentration of the deviations (consistent with a gradient-dominant Ξ rather than a resonance-like one), (ii) the consistency of deviation patterns across hadronic channels (disfavouring η_1-driven channel-dependent shifts), and (iii) the tight constraint on δC_7 from radiative modes (disfavouring strong η_3 coupling to $\mathcal{O}_7$). A first-principles derivation of the (η_1 : η_2 : η_3) hierarchy from the substrate is not currently available and is an open question. Stating this openly: the η_2-dominant regime is the phenomenological working hypothesis of this paper, not a derived result.

6.6 Summary of testable predictions

Table 2 consolidates the framework's concrete predictions, their current status, and the measurement programme that would decisively test them. The labels P1–P8 below are used independently of the §6.3 numbering (which shares "P1–P5" labels in a related but not identical context); Table 2 is the canonical reference.

Table 2. VERSF predictions for rare $b \to s\ell^+\ell^-$ observables, ordered by test timescale. "Current" refers to existing LHCb Run-2 data plus the 2022–2025 updates; "Upgrade II" refers to projected LHCb Upgrade-II precision (~2030+). Upgrade-II numerical targets are order-of-magnitude estimates based on expected luminosity scaling and should be verified against the LHCb physics-case documents at submission.

#	Prediction	Structural origin	Current status	How to test	Timescale	Falsifies if...
P1	LFU-universal δC_9 to $\mathcal{O}(m_\mu^2/m_b^2) \approx 5\times10^{-4}$	Hadronic-origin ΔS_V independent of lepton flavour	Passed at ~4% precision (R_K, R_K*, R_φ)	R_K, R_K*, R_φ at Upgrade-II (~1% precision)	Immediate–2030	R_K deviates from unity by > 0.1%
P2	$\delta C_{10}/\delta C_9$ controlled by λ_{10}/λ_9	Shared $(\partial\varphi)^2$ coupling of $\mathcal{C}_9$ and $\mathcal{C}_{10}$	Data prefers $\lambda_{10}/\lambda_9 \approx 0$; simplest guess $\lambda_{10}/\lambda_9 \approx 1$ excluded by $\mathscr{B}(B_s\to\mu\mu)$	Companion PGL(3,2) calculation	Immediate (theoretical)	Substrate calculation gives $|\lambda_{10}/\lambda_9| > 0.2$

#	Prediction	Structural origin	Current status	How to test	Timescale	Falsifies if...
P3	$\delta C_9(q^2)$ derivative-shaped in [1.1, 6] GeV^2	Gradient-dominant $\Delta S_V \Longrightarrow \Xi \propto \langle(\partial\varphi)^2\rangle(q^2)$	Untested (global fits assume q^2-constant δC_9)	Shape-unconstrained global fit of existing LHCb data	Immediate (reanalysis)	Best-fit $\delta C_9(q^2)$ is flat in q^2
P4	Null shift of $\mathcal{O}_9'$, $\mathcal{O}_{10}'$, $\mathcal{O}_S$, $\mathcal{O}_P$	$\delta C_i = -C_i^{SM} \cdot \lambda_i \cdot \Xi \Longrightarrow C_i^{SM} = 0$ gives zero shift	Passed at all current precisions (Table 1)	Upgrade-II angular analyses; $B_s \to \mu\mu$ precision improvement	2025–2035	Future measurements establish $\delta C_i'$ or $\delta C_{\{S,P\}}$ non-zero at $> 3\sigma$
P5	One-parameter locus across $B^+ \to K^+\mu\mu$, $B_s \to \varphi\mu\mu$, $B \to K^*\mu\mu$	Single $\Xi(q^2)$ weighted by channel-specific form factors	Untested (channels fit independently)	Constrained multi-channel global fit with $\delta C_i = \lambda_i \cdot \Xi(q^2)$	Immediate (reanalysis)	Channels prefer inconsistent Ξ shapes
P6	$\mathscr{B}(B \to K^{(*)}\tau^+\tau^-)$ enhanced by 20–30% at high q^2	Inherited from LFU-universal δC_9^U (**not VERSF-distinctive**)	Not yet measured	Upgrade-II and Belle II measurements	2030+	Null result at predicted magnitude
P7	Structural correction extends across $b \to s$ channels ($b \to s\mu\mu$, $B_s \to \mu\mu$, $b \to s\nu\bar{\nu}$)	Universality of pathway reweighting on the substrate	$b \to s\mu\mu$ and $B_s \to \mu\mu$ consistent; $b \to s\nu\bar{\nu}$ not yet measured	$\mathscr{B}(B \to K\nu\bar{\nu})$ at Belle II; high-precision $B_s \to \mu\mu$	2025–2030	$b \to s\nu\bar{\nu}$ deviation pattern incompatible with $b \to s\mu\mu$ framework
P8	$\lvert\lambda_7/\lambda_9\rvert \lesssim 0.7$ required by $\mathscr{B}(B \to X_s\gamma)$	Internal consistency of η_2-dominant regime	Constraint satisfied by η_2-dominance ansatz	Companion PGL(3,2) calculation	Immediate (theoretical)	Substrate calculation gives $\lvert\lambda_7/\lambda_9\rvert > 0.7$

The framework is at genuine empirical risk on multiple timescales: immediately through the companion PGL(3,2) calculation (P2, P8) and through reanalyses of existing LHCb data (P3, P5); over the current decade through Upgrade-II precision on LFU ratios and primed operators (P1, P4); and on a 2030+ horizon through $b \to s\tau\tau$ measurements (P6) and Belle-II $b \to s\nu\bar{\nu}$ data (P7). Failure of any of P1–P5 or P8 would falsify the leading-order framework; failure of P6 or P7 would require reconsideration but could be accommodated by higher-order structural corrections.

6.7 Application to $B_s \to \mu^+\mu^-$

The purely leptonic decay $B_s \to \mu^+\mu^-$ is an independent and particularly clean test of the framework's structural content, because its SM amplitude is dominated by $\mathcal{O}_{10}$ and its sensitivity to scalar, pseudoscalar, and primed operators is exquisite. The experimental value, averaged across LHCb, CMS, and ATLAS measurements, currently lies within ~10% of the SM prediction — a tight constraint on any new-physics contribution. §6.3.5 treated $B_s \to \mu\mu$ as a constraint on the pathway-ratio structure; here we treat it as a *second channel* in which the framework makes predictions, complementing the $b \to s\ell\ell$ analysis.

What the framework predicts. Under Theorem 3, the VERSF contribution to $B_s \to \mu^+\mu^-$ factorises through the same operator-graded pathway averaging as $b \to s\ell\ell$. Schematically,

$$\mathcal{B}(B_s \to \mu\mu)_{VERSF} / \mathcal{B}(B_s \to \mu\mu)_{SM} \approx |1 + \delta C_{10} / C_{10}^{SM}|^2 + \mathcal{R}\{S,P\},$$

where $\delta C_{10} = -C_{10}^{SM} \cdot \lambda_{10} \cdot \Xi(q^2 = m_{B_s}^2)$ and $\mathcal{R}_{\{S,P\}}$ encodes scalar and pseudoscalar contributions that vanish at leading order by Claim D. The full amplitude structure — including helicity-weighted mass-suppressed contributions from C_{10}', C_S, C_P and their interference — follows the standard form [18]; the expression above captures the leading structural content relevant here. Two features of the prediction are structurally determined and therefore constrain the framework non-trivially.

Cross-channel kinematic constraint. The shape function Ξ evaluated at $q^2 = m_{B_s}^2 \approx 28.9$ GeV² is the *same* $\Xi(q^2)$ that governs the $b \to s\ell\ell$ anomalies at $q^2 \in [1.1, 6]$ GeV². $B_s \to \mu\mu$ therefore tests the framework at a kinematic point far outside the anomaly window, with no additional free parameters. If the data-preferred shape of $\Xi(q^2)$ in the anomaly region is the derivative-localised profile of Claim C, its value at high $q^2 \approx m_{B_s}^2$ should be substantially smaller than its value at low q^2. Quantitatively, if $\Xi(28.9\ \text{GeV}^2) / \Xi(4\ \text{GeV}^2) \sim 0.1$–$0.3$ under a gradient-dominated regime (as expected from $\langle(\partial\varphi)^2\rangle$ being concentrated at low q^2), the VERSF contribution to $B_s \to \mu\mu$ is automatically suppressed by the same factor relative to the anomaly-window shift.

Axial-sector cross-check. The axial contribution is controlled by λ_{10}, not λ_9. The $B_s \to \mu\mu$ constraint therefore directly tests the pathway ratio λ_{10}/λ_9 discussed in §6.3 (P2) and §6.3.5. Working through the numbers: the simplest guess $\lambda_{10}/\lambda_9 \approx 1$, combined with $\lambda_9 \cdot \Xi(\text{anomaly window}) \approx 0.23$ (from $\delta C_9/C_9^{SM} \approx -0.23$ with $\delta C_9 \approx -1$), and assuming Ξ is not strongly suppressed at high q^2, would give $\delta C_{10} \approx +1$ and shift $\mathcal{B}(B_s \to \mu\mu)$ downward by approximately $|1 - 1/4.3|^2 \approx 0.59$, i.e. a ~41% deficit. (The precise number depends on the scheme-dependent value of C_{10}^{SM} at $\mu = m_b$; values in the range -4.1 to -4.3 give deficits between ~41% and ~43%.) The current measurement excludes such a deficit at several sigma. Under the data-preferred regime $|\lambda_{10}/\lambda_9| \lesssim 0.2$, the predicted $B_s \to \mu\mu$ shift is $\lesssim 10\%$ even without the high-q^2 suppression of Ξ, and is likely much smaller once the gradient-dominated Ξ shape is accounted for. The framework is therefore automatically consistent with current $B_s \to \mu\mu$ data in the same parameter regime that accommodates the $b \to s\ell\ell$ anomalies. This is a non-trivial cross-channel check: a framework tuned to reproduce $\delta C_9 \approx -1$ at low q^2 would generically produce a detectable effect at high q^2 through the same δC_{10}, and the framework passes this test only because of the independent structural features (gradient-dominated Ξ, axial-decoupled λ_{10}).

Claim D sharpness at B_s → μμ. The scalar and pseudoscalar sensitivity of B_s → μμ is the tightest known bound on δC_S and δC_P, roughly |δC_S|, |δC_P| ≲ 0.02 depending on the fit. Claim D predicts δC_S = δC_P = 0 at leading order. Consistency requires that next-to-leading-order VERSF corrections to these operators be below the same level, i.e. NLO terms in the Theorem 2 basis must contribute no more than ~0.5% of C_{10}^SM to the scalar/pseudoscalar sector. This is a constraint on the expansion in $(m_b/\Lambda_V)^2$ rather than on the leading-order framework itself, and is satisfied provided the NLO expansion converges at the expected rate.

Distinguishing VERSF from leptoquark scenarios at B_s → μμ. Leptoquark models that generate non-zero δC_S or δC_P (such as R_2) would produce a B_s → μμ signature distinct from the pure-δC_{10} pattern predicted here. A small scalar/pseudoscalar contribution changes the helicity structure of B_s → μμ, which has been proposed as a future discriminator through angular observables of the muons relative to the decay axis (though present statistics are insufficient for such measurements). VERSF predicts no such angular distortion at leading order. A future precision measurement of the B_s → μμ angular distribution, if it reveals helicity-structure deviations from the SM pure-$\mathcal{O}_{10}$ pattern, would falsify Claim D and distinguish VERSF from leptoquark alternatives; a null result would support both.

Upgrade-II sensitivity. Projected Upgrade-II precision on $\mathcal{B}$(B_s → μμ) is expected to reach ~4%, compared to ~10% currently. At that precision, the framework predicts no deviation from SM under any of its current regimes — a 4% null result would further confirm Claim D, while a deviation would require either NLO VERSF corrections above the expected scale or a breakdown of the gradient-dominated regime. Either outcome would be informative.

---

7. Epistemic status

Statement	Status
Pathway decomposition (P1)	**Postulated** (VERSF foundational axiom)
Boltzmann weighting, Theorem 0	**Derived** (from additivity, locality, closure)
Euclidean formulation with Wick rotation	**Assumed** (motivated by substrate structure; consistency not proved here)
SM recovery, Theorem 1	**Assumed at substrate level** (pathway-to-diagram map, CKM, GIM, RG running taken as inherited from substrate QFT reproduction; not independently derived here)
Structural correction basis, Theorem 2	**Derived** (given symmetry + derivative-count restriction)
Wilson-coefficient shift formula, Theorem 3	**Derived** (Euclidean linear response; magnitude of resulting δC_i carries ~25% NLO uncertainty from §4 truncation)

Statement	Status
Shared shape function $\Xi(q^2)$ at leading order (Claim A)	**Derived**
Null shift of vanishing-SM-coefficient operators (Claim D)	**Derived** (from Theorem 3 structure)
Claim B structural form (vector–axial linked via shared $(\partial\varphi)^2$)	**Derived**
Claim B numerical ratio λ_{10}/λ_9	**Open** (requires PGL(3,2) calculation; simplest guess $\lambda_{10}/\lambda_9 \approx 1$ disfavoured by data)
Operator ratios $\lambda_7 : \lambda_9 : \lambda_{10}$	**Open** (requires PGL(3,2) chirality calculation)
q^2-shape of Ξ (Claim C)	**Motivated and structurally expected, not yet computed** (qualitative argument only; $\langle(\partial\varphi)^2\rangle(q^2)$ not evaluated)
η_2-dominated regime	**Phenomenological working hypothesis** (selected by data; substrate derivation open)
Identification $\varphi \equiv \kappa$-field, $\Lambda\_V \equiv \xi^{-1}$	**Inherited** from upstream VERSF derivations
Leading-order EFT truncation at m_b scale	**Working approximation** with ~25% NLO uncertainty; full control requires NLO computation or $\Lambda\_V \gg 10$ GeV
Closure scale above b-mass ($\xi^{-1} \gtrsim 10$ GeV)	**Assumed**

---

8. Discussion

Relation to conventional new physics. A generic new-physics explanation introduces new ultraviolet degrees of freedom and then fits their low-energy imprint through independent Wilson-coefficient deformations. The present framework does the opposite. It does not add particles; it modifies the weighting of admissible transition histories at the substrate level, from which the Wilson shifts emerge only after coarse-graining. The resulting phenomenology is therefore more correlated than in a generic EFT fit: one structural profile, three operator sectors, fixed ratios in principle determined by pathway geometry rather than by separate fit parameters. The anomaly is not produced by a new particle in the loop, but by a structural reweighting of which loop-level transition histories are most admissible.

Falsifiability. The framework is falsified at leading order if any of the following occur:

- global fits require independent q^2-profiles for δC_7, δC_9, and δC_{10}, rather than a single shared structural profile $\Xi(q^2)$ (Claim A);
- the substrate-level PGL(3,2) calculation gives $\lambda_{10}/\lambda_9 \approx +1$ (the naïve equal-coupling ratio), which is currently in tension with $\delta C_{10} \approx 0$ preferred by global fits and excluded by $\mathscr{B}(B\_s \to \mu^+\mu^-)$ (Claim B);

- measurements establish non-zero shifts in the primed operators $\mathcal{O}_9'$, $\mathcal{O}_{10}'$ or in scalar/pseudoscalar operators at significance exceeding the Upgrade-II thresholds tabulated in §6.3.5, Table 1 (Claim D);
- the anomaly pattern is best described by a q^2-flat shift rather than a low-q^2-concentrated derivative-like profile (Claim C);
- future LFU measurements require a genuinely flavour-nonuniversal leading correction beyond the $\mathcal{O}(m_\ell^2/m_b^2)$ level, contrary to the hadronic-side structural origin of ΔS_V;
- the substrate-level calculation gives $|\lambda_7/\lambda_9| > 0.7$, which would produce δC_7 in conflict with the tight $\mathscr{B}(B \to X_s\gamma)$ constraint.

These conditions are mutually independent, providing multiple orthogonal routes to falsification.

Open calculations. The substrate-level derivation of the λ_i ratios from PGL(3,2) chirality assignments is the principal unfinished piece. A companion paper is indicated. Extension to $B_s \to \mu\mu$, $b\to s\nu\bar{\nu}$, and $\Lambda_b \to \Lambda\ell\ell$ provides further cross-channel tests: VERSF predicts the same $\Xi(q^2)$ controls all of them, with channel-specific λ computed from pathway geometry.

Limits of validity. The framework is an effective structural correction valid for $E \ll \Lambda_V$. It is not a UV completion of the SM. The UV completion is the VERSF substrate itself, which is discrete and has no continuum field-theoretic description above ξ^{-1}.

---

9. Conclusion

Rare $b\to s\ell^+\ell^-$ transitions provide a clean window onto the fine structure of admissible pathway weighting in VERSF. We have shown that:

1. The standard effective Hamiltonian is recovered in the limit of vanishing structural correction (Theorem 1);
2. There is a unique leading-order structural correction functional consistent with covariance, locality, and closure (Theorem 2);
3. The induced Wilson-coefficient shifts share a single q^2-shape function with operator-specific ratios fixed by pathway geometry (Theorem 3);
4. The post-2022 experimental pattern — LFU-preservation across R_K, R_{K^*}, R_φ combined with persistent C_9-dominant deviations at the 4σ–5σ level in low-q^2 branching fractions and the P_5' angular observable — is structurally consistent with — and partially predicted by — the framework. LFU-preservation is a direct consequence of ΔS_V being defined on hadronic-side fields that do not couple to lepton flavour; C_9-dominance is the natural consequence of the η_2 (gradient) term dominating; the low-q^2 localisation is structurally expected from the kinematic shape of $\langle(\partial\varphi)^2\rangle(q^2)$, subject to the shape calculation being completed.

Where conventional new-physics models originally targeted the LFU-violating picture and now struggle to accommodate its disappearance, VERSF accommodates the current pattern natively and without additional model-building.

The programme makes **fewer**, not more, free parameters than a generic new-physics interpretation, and supplies concrete falsifiable predictions. These span three timescales:

Already consistent with data:

- **Claim D, no primed or scalar/pseudoscalar shift**: consistent with current $B\_s \to \mu\mu$ and angular-fit bounds on $\mathcal{O}_9'$, $\mathcal{O}_{10}'$.

Immediately testable with existing LHCb data:

- **Claim A, shared shape**: a constrained multi-operator fit with $\delta C_i = \lambda_i \cdot \Xi(q^2)$ should perform comparably to independent $\delta C_i(q^2)$ fits.
- **Claim C, q^2-shape**: bin-by-bin analysis of δC_9 within [1.1, 6] GeV² should reveal the gradient-like kinematic profile indicated by $\langle(\partial\varphi)^2\rangle(q^2)$; a flat shift would disfavour the η_2-dominant regime.

Dependent on the companion PGL(3,2) calculation:

- **Claim B, vector–axial ratio**: whether the substrate-level λ_{10}/λ_9 lands near 0 (consistent with current fits) or near +1 (in tension with data) is the decisive test of the pathway-ratio structure.

Upgrade-II-era measurements of $B \to K^{(*)}\tau^+\tau^-$ will provide the next decisive test: a LFU-universal δC_9^U at the currently preferred magnitude predicts a 20–30% enhancement over SM, concentrated at high q^2. Observation at this level would further support the framework; a null result at Upgrade-II sensitivity would require reconsideration. In parallel, the purely leptonic channel $B\_s \to \mu^+\mu^-$ (§6.7) provides a high-q^2 cross-check on the same $\Xi(q^2)$ and λ_{10} structure governing the $b \to s\ell\ell$ anomalies at low q^2, and offers a distinct Upgrade-II sensitivity to Claim D via improved scalar/pseudoscalar bounds.

The principal unfinished step is now sharply identified: the first-principles derivation of the operator ratios $\lambda_7 : \lambda_9 : \lambda_{10}$ from the PGL(3,2) chirality structure of the K=7 substrate. This is not a technical detail, but the decisive bridge between a constrained phenomenological description and a parameter-free prediction.

If this step is completed successfully, rare B-meson decay anomalies cease to be merely signals of physics beyond the Standard Model, and instead become direct probes of the underlying closure structure of physical reality as described by VERSF.

This marks a shift in perspective: rare decay anomalies need not point to new degrees of freedom, but may instead be the first observable signatures of how physical amplitudes emerge from an underlying discrete commitment structure, making flavour anomalies a direct probe of the mechanism by which those amplitudes are constructed.

Appendix A — Dimensional audit

Convention: λ is the dimensionless tick-count parameter along Γ; the tick-to-proper-time mapping (each tick ~ Λ_V^{-1} in continuum) is absorbed into the η_n coefficients, so that every term displayed in ΔS_V is manifestly dimensionless without additional Λ_V factors in the integral measure.

Object	Dimension	Source
φ	0	Closure-field phase (assigned)
ρ_c, ρ_*	0	Normalised to ρ_*
$\Lambda_V = \xi^{-1}$	1 (mass)	Inverse closure scale
λ (tick count)	0	Dimensionless tick parameter on K=7
$d\lambda$	0	Dimensionless differential
$\partial_\mu\varphi$	1	Spacetime gradient of dim-0 field (continuum limit)
$(\partial\varphi)^2$	2	Contracted gradient
$\Box\varphi$	2	Laplacian of dim-0 field
$(\partial\varphi)^2/\Lambda_V^2$	0	Dimensionless
$\Box\varphi/\Lambda_V^2$	0	Dimensionless
$\delta\rho_c/\rho_*$	0	Dimensionless
η_n	0	Dimensionless (absorbs tick-to-proper-time map)
$\int_\Gamma d\lambda\,(\cdot)$	inherits dim of $(\cdot)$	Tick sum (not a length-weighted integral); no dimension is added by the measure
Each term in ΔS_V	0	✓ (integrand dimensionless by construction)
ΔS_V	0	✓

Appendix B — Relation to upstream VERSF papers

- **Closure scale ξ**: derived from ledger completeness, finite distinguishability, and compact cyclic closure on K=7 (Taylor, *Closure Scale ξ from Three Axioms*) [12].
- **κ-field mass $m_\kappa^2 = (3/4)\,\xi^{-2}$**: Projection Theorem + PGL(3,2) irreducibility (Taylor, *κ-Field Mass Derivation*) [13].
- **Single-Source Theorem**: $\rho(x,t)$ as unique source of physical observables (Taylor, *Single-Source Theorem*) [14].
- **K=7 commitment architecture and PGL(3,2) structure**: foundational substrate results (Taylor, VERSF programme papers) [15].
- **Spectral density $J(\omega)$**: Caldeira-Leggett formalism applied to K=7 commitment bath (Taylor) [16].

The present paper uses ξ, m_κ, and ρ as fixed inputs; it does not re-derive them.

References

Standard framework

[1] G. Buchalla, A. J. Buras, and M. E. Lautenbacher, "Weak decays beyond leading logarithms," *Rev. Mod. Phys.* **68**, 1125 (1996). Standard reference for the $\Delta B = 1$ effective Hamiltonian and operator basis $\{\mathcal{O}_7, \mathcal{O}_9, \mathcal{O}_{10}\}$.

Experimental measurements

[2] LHCb Collaboration, R. Aaij *et al.*, "Test of lepton universality in $b \to s\ell^+\ell^-$ decays," *Phys. Rev. Lett.* **131**, 051803 (2023); "Measurement of lepton universality parameters in $B^+ \to K^+\ell^+\ell^-$ and $B^0 \to K^{*0}\ell^+\ell^-$ decays," *Phys. Rev. D* **108**, 032002 (2023). The 2022 R_K, R_K* update returning values consistent with unity.

[3] LHCb Collaboration, R. Aaij *et al.*, "Test of lepton flavour universality with $B_s^0 \to \varphi\ell^+\ell^-$ decays," arXiv:2410.13748 [hep-ex] (2024). R_φ measurement in three q^2 bins, consistent with SM.

[4] LHCb Collaboration, R. Aaij *et al.*, "Measurement of the branching fraction ratio R_K at large dilepton invariant mass," arXiv:2505.03483 [hep-ex] (2025). $R_K = 1.08\,^{+0.11}_{-0.09}$ (stat) ± 0.04 (syst) above the $\psi(2S)$ mass.

[5] T. Hurth, F. Mahmoudi, and S. Neshatpour, "Data-driven analyses and model-independent fits for present $b \to s\,\ell\ell$ results," arXiv:2508.09986 [hep-ph] (2025, v2 Nov 2025). Updated global fit incorporating the 2025 LHCb angular analysis and 2024 CMS data; finds deviations in $\mathcal{B}(B^+ \to K^+\ell^+\ell^-)$ at 4σ–5σ in the low-q^2 region.

[6] LHCb Collaboration, R. Aaij *et al.*, "Branching Fraction Measurements of the Rare $B_s^0 \to \varphi\mu^+\mu^-$ and $B_s^0 \to f_2'(1525)\mu^+\mu^-$ Decays," *Phys. Rev. Lett.* **127**, 151801 (2021). The 3.6σ deficit in $\mathcal{B}(B_s^0 \to \varphi\mu^+\mu^-)$ for $q^2 \in [1.1, 6]$ GeV².

[7] LHCb Collaboration, R. Aaij *et al.*, "Measurement of CP-averaged observables in the $B^0 \to K^{*0}\mu^+\mu^-$ decay," *Phys. Rev. Lett.* **125**, 011802 (2020). The LHCb angular analysis establishing the persistent P_5' tension.

[8] S. Descotes-Genon, T. Hurth, J. Matias, and J. Virto, "Optimizing the basis of $B \to K^*\,\ell^+\ell^-$ observables in the full kinematic range," *JHEP* **05** (2013) 137. The original definition of P_5' as a theoretically clean angular observable.

[9] CMS Collaboration, "Angular analysis of the decay $B^0 \to K^{*0}\mu^+\mu^-$ from proton-proton collisions at $\sqrt{s} = 13$ TeV," CMS-BPH-21-002 (2024). Independent confirmation of the LHCb angular-observable tensions at the 1σ level.

Global fits and theoretical analyses

[10] M. I. Ali *et al.*, "Constraints on lepton flavor universal and non-universal New Physics in $b \to s\ell^+\ell^-$ decays: a global SMEFT survey," arXiv:2502.20145 [hep-ph] (2025). SMEFT global fit finding P_5' deviation at 3.3σ and $\mathcal{B}(B_s \to \varphi\mu^+\mu^-)$ at 3.6σ, with best-fit LFU-universal δC_9^U scenarios.

[11] M. Algueró, B. Capdevila, S. Descotes-Genon, J. Matias, and M. Novoa-Brunet, "$b \to s\ell^+\ell^-$ global fits after Moriond 2023 results," *Eur. Phys. J. C* **83**, 648 (2023). Global-fit analysis establishing the preference for LFU-universal contributions to C_9 following the 2022 R_K update.

Leptoquark phenomenology

[17] I. Doršner, S. Fajfer, A. Greljo, J. F. Kamenik, and N. Košnik, "Physics of leptoquarks in precision experiments and at particle colliders," *Phys. Rep.* **641** (2016) 1–68, arXiv:1603.04993 [hep-ph]. Comprehensive review of leptoquark phenomenology, used for the comparison of VERSF null-shift predictions with S_3, R_2, and U_1 leptoquark scenarios in §6.3.5.

$B_s \to \mu\mu$ amplitude structure

[18] C. Bobeth, M. Gorbahn, T. Hermann, M. Misiak, E. Stamou, and M. Steinhauser, "$B_{s,d} \to \ell^+\ell^-$ in the Standard Model with reduced theoretical uncertainty," *Phys. Rev. Lett.* **112** (2014) 101801, arXiv:1311.0903 [hep-ph]. Standard reference for the SM $B_s \to \ell^+\ell^-$ amplitude and its beyond-SM operator structure, including scalar/pseudoscalar contributions and helicity-weighted corrections; used in §6.7 as the reference form behind the schematic expression. *Note: the exact reference combination for the BSM $B_s \to \mu\mu$ amplitude structure should be verified at submission against the conventions of the chosen global-fit paper.*

Upstream VERSF programme

[12] K. Taylor, "The Closure Scale ξ from Three Axioms," VERSF programme paper (AIDA Institute, versf-eos.com).

[13] K. Taylor, "The κ-Field Mass Derivation via the Projection Theorem and PGL(3,2) Irreducibility," VERSF programme paper.

[14] K. Taylor, "The Single-Source Theorem: ρ(x,t) as the Source of Physical Observables," VERSF programme paper.

[15] K. Taylor, "K=7 Commitment Architecture and PGL(3,2) Structure," VERSF programme foundational series.

[16] K. Taylor, "Spectral Density J(ω) of the Commitment-Event Bath from the K=7 Architecture," VERSF programme paper.