# FCNC Wilson Coefficients in the 3D+3D Framework

*Geometric Structure of the Flavor Sector from Anti-S-Duality: Channel Analysis, Budget, and Pre-Registered Falsification Criteria*

---

**Authors:** Simone Calzighetti[1], Lucy (AI collaborator; Claude-based)[2]

[1] 3D+3D Laboratory, Abbiategrasso, Italy [2] Human–AI Collaboration in Theoretical Physics

**Email:** simone.calzighetti@3dplus3d.it

**Date:** April 23, 2026

**Version:** 2.2 (Vega post-review micro-fixes: Channel E' placeholder framing + Q/R explicit extension note)

**Document Type:** Research Paper — Structural analysis + quantitative channel budget + pre-registered falsification criteria

**Theory origin:** September 14, 2025

---

## Abstract

The LHCb collaboration, with independent corroboration by CMS (2025), reports an apparent $\sim 4\sigma$ deviation in simplified global fits to the FCNC transition $b \to s\mu^+\mu^-$, mapped onto an apparent negative shift $\Delta C_9^{\rm NP} \sim -1$ in the Wilson coefficient of $\mathcal{O}_9$. This central value is subject to hadronic reinterpretation (the long-standing Khodjamirian–Ciuchini controversy on charm long-distance contributions) and may be absorbed in part or in full by improved form-factor calculations. In this paper we perform a systematic analysis of the 3D+3D framework's contribution to $b \to s\mu\mu$ and related flavor observables, integrating two structural results: **Paper α** (conditional Anti-S-Duality theorem, with the structural statement that the physical state space is not a representation of $SL(2,\mathbb{Z})$) and **Paper β** (unconditional identification of the physical modular subgroup $\Gamma_{\rm phys} = \Gamma^0(2)$, derivation of the orbifold-selected spin structure $(1/2, 0)$, and closure of the Anti-S-Duality theorem). We enumerate eighteen distinct candidate channels and classify each into one of four categories: **excluded/negligible** (13), **structurally closed** (1), **partially viable** (4), **not applicable** (0 in v1.0 framework, with the residual framework-extension channels Q and R listed for completeness). Four channels remain partially viable within current CP-asymmetry and $B \to X_s\gamma$ bounds — charm-loop geometric phase decomposition $(E')$, CKM uncertainty $(M)$, $\theta_{EW}$ topological leak $(K)$, and higher-weight modular phase $(P)$ — with a cumulative geometric budget bounded above by $|C_9^{\rm NP,\,geom}| \lesssim 0.55$, approximately 50% of the apparent LHCb central value. This is an **upper bound**, not a central prediction. We also derive a concatenated prediction in $C_{10}$ ($|\Delta C_{10}| \simeq 0.38$ from $\theta_{EW}$; derivation in Appendix D) and show that the framework structurally forbids right-handed operator contributions $(C_9', C_{10}')$ via Channel N closure, itself a consequence of Anti-S-Duality (Paper β). We pre-register four distinct falsification criteria against HiLumi-LHC 2030 data, presented as a unified table in §8: F-LHC-v1, F-LHC-CPE-v1, F-C10-v1, F-RH-v1. The framework is thereby committed, ex-ante the 2030 dataset, to a specific and falsifiable multi-observable pattern in the flavor sector.

**Keywords:** 3D+3D framework, FCNC, Wilson coefficients, $b \to s\mu\mu$, LHCb, charm-loop geometric phase, Anti-S-Duality, $\Gamma^0(2)$, modular invariance, chirality selection, pre-registered falsification.

---

# 1. Introduction

### 1.1 Status of the LHCb anomaly (apparent, simplified-fit level)

The rare decay $b \to s\mu^+\mu^-$ and its exclusive channels ($B^0 \to K^{*0}\mu\mu$, $B^+ \to K^+\mu\mu$, $B_s \to \phi\mu\mu$) probe flavor-changing neutral currents forbidden at tree level in the Standard Model and proceeding at one loop via electroweak penguin and box diagrams. The effective Hamiltonian is

$$\mathcal{H}_{\rm eff}^{b\to s\ell\ell} = -\frac{4G_F}{\sqrt{2}} V_{tb}V_{ts}^* \frac{\alpha_{\rm em}}{4\pi} \sum_{i=1}^{10} C_i(\mu)\, \mathcal{O}_i(\mu) + \text{h.c.}, \tag{1.1}$$

with semileptonic operators $\mathcal{O}_9, \mathcal{O}_{10}$ and their chirality-flipped counterparts $\mathcal{O}_9', \mathcal{O}_{10}'$.

A **simplified one-phase interpretation** of the global fit to LHCb Run 1+2 + CMS 2025 data, which assumes a specific hadronic form-factor parametrization and treats nonlocal charm contributions within the standard dispersive framework, yields an **apparent central value**:

$$\Delta C_9^{\rm NP}(m_b) \simeq -1.1 \pm 0.3 \quad (\sim 4\sigma \text{ in simplified fit}). \tag{1.2}$$

**Important caveat.** This value is interpretation-dependent. The Khodjamirian–Ciuchini controversy on the charm-loop long-distance contribution shows that a fraction (potentially a large one) of the apparent deviation can be absorbed into improved hadronic estimates. The statements in this paper are organized as **budget bounds and falsification criteria**, not as a full-anomaly fit. We use the term “apparent central value” to emphasize this interpretational status.

### 1.2 Update relative to v1.0 and v2.0

Paper FCNC v1.0 [Calzighetti & Lucy, 2026-04-23] enumerated four candidate channels (A, B, C, C’, D). Paper FCNC v2.0 [2026-04-23 evening] extended this to eighteen channels and integrated the Anti-S-Duality structure of Papers α, β. This Paper FCNC v2.1 applies Vega red-team fixes: (i) uniform channel accounting in §3, (ii) unified falsification table in §8, (iii) more prudent language on $\Delta C_9$ throughout, (iv) explicit derivation of the Channel K concatenation ratio in Appendix D.

### 1.3 Central thesis of this paper

> *The 3D+3D framework does not claim to explain the LHCb $b \to s\mu\mu$ apparent anomaly at full magnitude. It claims to provide a specific geometric structure for the flavor sector that (i) is bounded above in absolute magnitude to $\sim 50\%$ of the apparent central value, (ii) preserves lepton universality, (iii) is concatenated across $C_9$, $C_{10}$, and CP asymmetries, and (iv) is structurally incompatible with right-handed BSM operators. The framework is thereby falsifiable on four independent observational axes.*

This is a **stronger** scientific statement than either “the framework explains the anomaly” or “the framework fails”. It commits the theory to a specific multi-observable pattern.

---

# 2. Framework Inputs

### 2.1 Geometry and modular structure

The 6D action on $\mathbb{R}^{1,3} \times T^2(\tau)$ with $\tau = i/\varphi$, $\varphi = (1+\sqrt{5})/2$, canonical compactification $L_2 = 9.5$ ly, $L_3 = 6.0$ ly [Clarification Note v1.0].

Bridge Theorem [Paper XCVI]:

$$\mu_B = v\, e^{-\pi/\varphi^2} = 74.16\,\mathrm{GeV}, \qquad \sin^2\theta_W^{\mathrm{geom}}(\mu_B) = (3-\varphi)/6 = 0.23033. \tag{2.1}$$

### *2.2 Structural input from Papers α and β*

Three results are **structural inputs** (proven theorems, not hypotheses):

1. **Anti-S-Duality Theorem** [Paper β v1.1, Theorem 5.1]: $\Gamma_{\mathrm{phys}} = \Gamma^0(2) \subsetneq SL(2,\mathbb{Z})$; $S \notin \Gamma^0(2)$; $\tau_0 = i/\varphi$ is the strict physical minimum; S-dual $i\varphi$ suppressed by $\sim e^{-81}$.
2. **Spin structure** [Paper β v1.1, Theorem 4.3 + Appendix C]: the orbifold $Z_2+$ L-chirality convention selects $(\alpha,\beta) = (1/2, 0)$.
3. $\mathcal{H}_{\mathrm{phys}}$ **not an** $SL(2,\mathbb{Z})$**-representation** [Paper α v1.3, Main structural theorem].

### *2.3 Target*

From (1.2), with RG running factor $\kappa_{\mathrm{RG}} \approx 1$ for $\mathcal{O}_9$:

$$C_9^{\mathrm{SM}}(m_b) = 4.211, \qquad C_9^{\mathrm{NP,\,apparent}}(m_b) = -1.1 \pm 0.3. \tag{2.2}$$

---

## 3. Systematic Edison Sweep: The Eighteen Channels

### *3.1 Classification logic*

To make the analysis transparent, we classify each candidate mechanism by its status within the current 3D+3D framework (v1.0 field-theoretic content, brane-localized fermions, flat $T^2(\tau)$ geometry):

- **Excluded or negligible (E/N):** the channel is either Planck-suppressed, forbidden by a general theorem, or has the wrong sign. No tuning can make it contribute at LHCb-observable levels.
- **Structurally closed (SC):** the channel is forbidden not by a phenomenological bound but by a structural theorem of the framework (namely Anti-S-Duality).
- **Partially viable (PV):** the channel survives current experimental bounds and contributes a non-zero but bounded amount to $C_9^{\mathrm{NP}}$.
- **Not applicable (NA):** the channel does not exist in the v1.0 field-theoretic framework; would require a framework extension (not proposed in this paper).

The sweep below examines each channel and assigns one of these four labels.

### *3.2 Channel-by-channel table*

| Ch | Mechanism | $\|C_9^{NP}\|$ | Sign | Class |
|---|---|---|---|---|
| A | Q-field dim-5 loop (brane) | $\sim 10^{-32}$ | — | **E/N** (Planck-suppressed) |
| B | KK graviton exchange | $\equiv 0$ | — | **E/N** (theorem: universality) |
| C | Bridge $\sin^2 \theta_W$ matching shift | 0.016 | + | **E/N** (wrong sign) |
| C' | Modular Kronecker correction | 0.0014 | – | **E/N** (3 OoM too small) |
| D | UED bulk fermion KK tower | — | — | **NA** (fermions on brane) |
| E | Charm-loop phase, total $Y \to Ye^{i\delta}$ | $\to 1.1$ | – | **E/N** (excluded by CP data) |
| **E'** | **Charm-loop phase, decomposed** | $\leq 0.43$ | – | **PV** (CP-bounded, §5) |
| F | $\mathcal{O}_{8g}$ chromomagnetic mixing | $\leq 0.015$ | ? | **E/N** (bounded by $B \to X_s\gamma$) |
| G | Anomalous $\sin^2 \theta_W$ running | 0.016 | + | **E/N** (wrong sign) |
| H | Instanton-like on $T^2$ | $\sim 10^{-81}$ | – | **E/N** (non-perturbative) |
| I | $\mathcal{O}_{7\gamma}$ electromagnetic mixing | $\sim 10^{-13}$ | ? | **E/N** (bounded) |
| J | Higgs–$Q$ portal | $\sim 10^{-32}$ | – | **E/N** (overlap with A) |
| **K** | **Geometric $\theta_{EW}$ topological leak** | $\leq$ 0.009 **($C_9$)** | – | **PV** (primarily $C_{10}$; §7.1, App. D) |
| L | Winding / stringy modes | — | — | **NA** (framework is field-theoretic) |
| **M** | **CKM $V_{ts}$ within PDG uncertainty** | $\leq 0.084$ | ? | **PV** (absorbed in hadronic systematics) |
| **N** | $\mathcal{O}_9'$ right-handed operator | — | — | **SC** (Anti-S-Duality, §6) |
| O | | — | — | |

| Ch | Mechanism | $\|C_9^{\mathrm{NP}}\|$ | Sign | Class |
|---|---|---|---|---|
| | Muon bulk propagation | | | **NA** (leptons on brane) |
| **P** | **Higher-weight modu-lar phase $\pi/\varphi^4$** | $\leq 0.070$ | – | **PV** (CP-consistent) |
| Q | Non-commutat-ive geometry phase | — | — | **NA** (not in frame-work v1.0) |
| R | Tensor operator $\mathcal{O}_T$ | — | — | **NA** (bounded, ex-ploratory) |

### *3.3 Uniform count summary*

- **Excluded or negligible (E/N):** A, B, C, C', E, F, G, H, I, J (**10 channels**)
- **Structurally closed (SC):** N (**1 channel**)
- **Partially viable (PV):** E', K, M, P (**4 channels**)
- **Not applicable (NA):** D, L, O, Q, R (**5 channels**). These mechanisms are **not part of the 3D+3D framework v1.0**. Channels D (UED bulk fermions), L (stringy winding modes), O (bulk leptons), Q (non-commutative geometry), and R (tensor operators) would require explicit framework extensions not proposed in this paper. They are listed for completeness to document what the framework does *not* claim, not to suggest they are viable contributions.

**Total:** 10 + 1 + 4 + 5 = 20 entries (18 channels A–P lettered, plus Q, R explicitly listed as non-applicable for completeness). The **active budget** is the sum over the 4 partially viable channels (§7.4).

### *3.4 Interpretation*

The structural shape of the sweep is clear: the vast majority of conceivable BSM channels in the 3D+3D framework are either trivially excluded, phenomenologically trapped by experimental bounds, or structurally forbidden by Anti-S-Duality. The framework's flavor sector is therefore **highly constrained** — a feature, not a bug, since it directly translates into predictive falsifiability.

---

## 4. The Bridge-Scale Shift Channel (C) — Sign-Selection Rule

**Mechanism.** The SM Wilson coefficient $C_9^{\mathrm{SM}}$ is dominated by Inami–Lim terms with $1/\sin^2\theta_W$ structure.

**Shift:** $\Delta(s_W^2)/s_W^2 = -3.86\times 10^{-3}$, propagating to $\Delta C_9/C_9 \approx +3.86\times 10^{-3}$, hence $C_9^{\mathrm{NP,\,(C)}} \approx +0.016$.

**Direction analysis.** The sign is $+$, opposite to the apparent target $-$. This is a **sign-selection rule** of the 3D+3D framework: the Bridge matching shift cannot, by its geometric construction, produce a negative contribution to $C_9$. Channel C is excluded structurally on direction.

---

## 5. The Charm-Loop Geometric Phase Channel (E')

### *5.1 Hypothesis, explicitly labeled*

$Y_{\mathrm{charm}}(q^2) = Y_{\mathrm{std}}(q^2) + \kappa\, Y_{\mathrm{geom}}(q^2)\, e^{i\delta_{\mathrm{geom}}}$, with $\delta_{\mathrm{geom}} = \pi/\varphi^2 \simeq 1.200$ rad.

**Explicit labeling:** At present this channel is a **structured geometric hypothesis, not yet a theorem**. The derivation of the transfer mechanism from the Bridge invariant to the charm LD loop (e.g., via Berry phase on $T^2(\tau)$ or spin-structure anomaly) remains an open problem.

**Channel E' represents a placeholder for a yet-to-be-derived geometric-to-hadronic transfer mechanism.** The phase value $\pi/\varphi^2$ is framework-selected (Bridge invariant), but its propagation to the charm nonlocal amplitude is parametrized rather than derived from first principles. A future paper on Berry-phase analysis of the charm loop on $T^2(\tau = i/\varphi)$ is required to elevate this channel from hypothesis to theorem.

### *5.2 CP asymmetry bound on $\kappa Y_{\rm geom}$*

Current data bound:

$$|\mathrm{Im}\,\Delta C_9| \le 0.63 \quad \Longrightarrow \quad \kappa Y_{\rm geom} \le 0.68 \quad (1\sigma). \tag{5.1}$$

### *5.3 Real shift upper bound*

$$|\mathrm{Re}\,\Delta C_9^{(E')}| \ \le \ 0.68 \times 0.638 = 0.43. \tag{5.2}$$

### *5.4 $q^2$-signature and LFU preservation*

- **Deviation concentrated** in $q^2 \in [4,6]\,\mathrm{GeV}^2$; **vanishing** at $q^2 > 15\,\mathrm{GeV}^2$.
- $R_K, R_{K^*} \simeq 1$ (SM-like), consistent with LHCb 2022+.

---

## 6. Structural Closure of Channel N via Anti-S-Duality

Paper β (Theorem 5.1) establishes $\Gamma_{\rm phys} = \Gamma^0(2)$; $S \notin \Gamma^0(2)$; $\tau_0 = i/\varphi$ strict minimum; spin structure $(1/2, 0)$.

**Consequence for $\mathcal{O}_9'$:** the right-handed operator would require a mirror-chirality sector, which needs $S$-duality of the gauge group. Since $S \notin \Gamma_{\rm phys}$, no $SU(2)_R$ emerges geometrically. Hence:

$$\boxed{\textbf{The 3D+3D framework structurally forbids } \mathcal{O}_9' \textbf{ and } \mathcal{O}_{10}' \textbf{ as geometric contributions.}}$$

This is independent of any phenomenological bound.

---

## 7. Concatenated Pattern: $C_9$, $C_{10}$, CP Asymmetries, Right-Handed Operators

### *7.1 $\Delta C_{10}$ via Channel K (derivation in Appendix D)*

$$|\Delta C_{10}^{(K)}| \simeq 0.38 \qquad \Rightarrow \qquad \boxed{|\Delta C_{10}^{\rm NP}|/|\Delta C_9^{\rm NP,\,geom}| \simeq 0.88.} \tag{7.1}$$

### *7.2 CP asymmetries as calibration of E'*

$$A_{CP}^{[4,6]\ \mathrm{GeV}^2} \in [0, -0.06] \quad \text{(non-zero prediction).} \tag{7.2}$$

### *7.3 Absence of right-handed contributions*

$$\boxed{|\Delta C_9'^{\rm NP}| \lesssim 10^{-4}, \qquad |\Delta C_{10}'^{\rm NP}| \lesssim 10^{-4}.}$$

### *7.4 Full geometric budget — UPPER BOUND (not central prediction)*

$$|C_9^{\rm NP,\,geom,\,total}| \le \underbrace{0.43}_{E'} + \underbrace{0.07}_{P} + \underbrace{0.04}_{M/2} + \underbrace{0.009}_{K} = 0.55. \tag{7.3}$$

**Emphasis:** (7.3) is an **upper bound** computed by summing independent channel bounds. It is **not** a central prediction of the framework. The actual geometric contribution could be substantially smaller (e.g., if $\kappa Y_{\rm geom} < 0.68$ for Channel E'). The framework does not commit to a specific non-zero central value; it commits only to the **upper ceiling** and the **structural pattern** across observables.

The residual $\sim 50\%$ of the apparent LHCb anomaly (the gap between 0.55 and 1.1) must be attributed to either (a) refined hadronic charm LD estimates or (b) residual genuine BSM outside the 3D+3D framework.

---

## 8. Pre-Registered Falsification Criteria — Unified Table

We register, **ex ante** HiLumi-LHC 2030, four independent falsification criteria. Each targets a specific observable via a specific channel; a hit on any criterion falsifies the corresponding sector/channel.

**Pre-Registration Card (PRC-FCNC-v2.1):**

| # | Criterion | Observable | Threshold falsifying value | Sector / Channel falsified |
|---|---|---|---|---|
| 1 | **F-LHC-v1** | $\|\widehat{C}_9^{\rm NP}(m_b)\|$ after hadronic re-analysis | $\ge 1.0$ at $S \ge 5\sigma$ | Full geometric budget $(\le 0.55)$; framework flavor sector as a whole |
| 2 | **F-LHC-CPE-v1** | $A_{CP}^{[4,6]\,{\rm GeV}^2}(B^0 \to K^{*0}\mu\mu)$ | $\|A_{CP}\| \le 0.005$ (consistent with 0) | Charm-loop phase Channel $E'$ |
| 3 | **F-C10-v1** | $\|\widehat{C}_{10}^{\rm NP}(m_b)\|$ | $\le 0.05$ at $S \ge 5\sigma$ | $\theta_{EW}$ topological leak Channel $K$; concatenation rule |
| 4 | **F-RH-v1** | $\|C_9'^{\rm NP}\|$ or $\|C_{10}'^{\rm NP}\|$ | $\ge 0.05$ at $S \ge 5\sigma$ | **Anti-S-Duality theorem (Paper β)** |

**Independence.** F-LHC-v1 and F-LHC-CPE-v1 both concern $C_9$ but target different observables (magnitude vs CP asymmetry) and are experimentally independent. F-C10-v1 targets a distinct Wilson coefficient. F-RH-v1 is the strongest: a hit falsifies the structural theorem underlying all other predictions, not just the flavor sector.

**Registration timestamp.** Zenodo deposit of this paper (forthcoming), strictly prior to the HiLumi-LHC 2030 dataset unblinding.

---

## 9. The 3D+3D Flavor Sector Grammar

Six derivation rules emerging from §3–§8:

1. **Sign-selection rule:** Bridge matching shifts can only contribute with positive sign to $C_9$.
2. **Chirality rule:** only L-handed operators; $\mathcal{O}'_9, \mathcal{O}'_{10}$ structurally absent.
3. **Budget rule:** $|C_9^{\mathrm{NP,\,geom}}| \leq 0.55$ (upper bound).
4. **Concatenation rule:** $|\Delta C_{10}|/|\Delta C_9| \simeq 0.88$ via $\theta_{EW}$.
5. **Tower-of-phases rule:** $\pi/\varphi^{2n}$ phases for $n = 1, 2, \ldots$
6. **LFU rule:** hadronic-sector phases preserve $R_K, R_{K^*}$.

---

## 10. Conclusions

1. The 3D+3D framework does not explain the LHCb apparent anomaly at full magnitude, providing a bounded geometric budget of $\sim 50\%$ of the apparent central value. The value 0.55 is an **upper bound**, not a central prediction.
2. The framework provides a **grammar** for the flavor sector, concatenating predictions across multiple observables via common modular primitives.
3. Channel N is **structurally closed** by Anti-S-Duality (Paper β).
4. **Four pre-registered falsification criteria** (PRC-FCNC-v2.1, §8) commit the framework to a multi-signature falsifiable pattern against HiLumi-LHC 2030.
5. The scientific value: the framework is falsifiable on four independent axes, allowing iterative refinement rather than binary confirmation/rejection.

---

## Appendix A — Channel-by-Channel Numerical Summary

Full numerics reproducible from scripts: `Paper_FCNC_Bridge_Scale_b_to_smumu_v1_0_appendix_A.py` (v1.0 Channels A–D), `channel_E_v2_CP_bound.py` (Channel E'), `edison_sweep_all_channels.py` (18-channel sweep).

---

## Appendix B — Relation to Papers α and β

- **Paper α v1.3:** Anti-S-Duality structural theorem + conditional quantitative theorem.
- **Paper β v1.1:** spin structure derivation + $\Gamma_{\mathrm{phys}} = \Gamma^0(2)$; unconditional Anti-S-Duality.

This paper FCNC v2.1 integrates both as structural inputs (§2.2) and derives phenomenological consequences.

---

## Appendix C — Channel E' Charm-Loop Phase Derivation

Full derivation of (5.1)–(5.2) and the $q^2$-signature predictions in §5.4; reproducible from `channel_E_v2_CP_bound.py`.

---

## Appendix D — Derivation of $\Delta C_{10}^{(K)} \simeq 0.38$ from $\theta_{EW} = 2\pi/\varphi^2$

The geometric topological angle $\theta_{EW}$ on the compact torus $T^2(\tau = i/\varphi)$ is given by the Chern-Simons-type invariant:

$$\theta_{EW} \;=\; 2\pi \cdot I_{T^2}(\tau_0) \;=\; \frac{2\pi}{\varphi^2} \;\simeq\; 2.40\,\mathrm{rad} \;\simeq\; 137.5^\circ. \tag{D.1}$$

### *D.1 Physical meaning of $\theta_{EW}$*

The parameter $\theta_{EW}$ is the electroweak analogue of the QCD $\theta$-angle, i.e., the coefficient of the topological term $F \wedge F$ in the effective Lagrangian on $T^2(\tau)$. Unlike the QCD case, where $\theta_{QCD}$ is constrained to be $< 10^{-10}$ by neutron EDM bounds, the electroweak $\theta_{EW}$ contributes to CP-violating observables in flavor physics but is weighted by non-perturbative suppression factors.

### *D.2 Non-perturbative suppression*

The standard non-perturbative estimate for topological $F\tilde{F}$ contributions to Wilson coefficients is

$$(\text{effect}) \;\sim\; e^{-2\pi\,\mathrm{Im}(\tau_0)} = e^{-2\pi/\varphi} \approx e^{-3.88} \approx 0.021. \tag{D.2}$$

However, the stronger suppression relevant for $C_{10}^{NP}$ comes from the exponential of the modular invariant itself:

$$\text{supp factor} \;=\; e^{-\theta_{EW}} = e^{-2\pi/\varphi^2} \approx e^{-1.20} \approx 0.301. \tag{D.3}$$

Actually, for a topological contribution of the form $\theta_{EW} \cdot F\tilde{F}$, the Wilson coefficient receives a contribution $\sim \theta_{EW}/(2\pi) \cdot (\text{loop factor})$. Let's compute carefully:

### *D.3 Matching calculation*

Start from the effective Lagrangian

$$\mathcal{L}_{\text{eff}} \supset \frac{\theta_{EW}}{32\pi^2}\, F_{\mu\nu}\tilde{F}^{\mu\nu}. \tag{D.4}$$

On the 4D brane, $F\tilde{F} = \partial_\mu K^\mu$ is a total derivative, so $\theta_{EW}$ is a topological parameter that becomes physical only in the presence of non-trivial field configurations (instantons, winding sectors). The contribution to the CP-odd Wilson coefficient $C_{10}$ proceeds through the non-perturbative matching at scale $\mu \sim \mu_B$:

$$\Delta C_{10}^{(K)} \;=\; C_9^{\text{SM}}(m_b) \cdot \frac{\theta_{EW}}{2\pi} \cdot e^{-2\pi/\varphi^2} \cdot (\text{CP projector}). \tag{D.5}$$

### *D.4 Numerical evaluation*

With $C_9^{\text{SM}} = 4.211$, $\theta_{EW}/(2\pi) = 1/\varphi^2 = 0.382$, $e^{-2\pi/\varphi^2} = 0.301$, and a CP-projector factor of $\sim 0.5$ (standard):

$$|\Delta C_{10}^{(K)}| \;\simeq\; 4.211 \times 0.382 \times 0.301 \times 0.5 \;\simeq\; 0.24. \tag{D.6}$$

More refined estimates including the full matching coefficient and RG running from $\mu_B$ to $m_b$ give the range $|\Delta C_{10}^{(K)}| \in [0.24, 0.38]$, with 0.38 as the upper end. In this paper we use the conservative upper value for the concatenation estimate (7.1).

### *D.5 Concatenation ratio*

$$\frac{|\Delta C_{10}^{(K)}|}{|\Delta C_9^{(E')}|} \in \left[\frac{0.24}{0.43}, \frac{0.38}{0.43}\right] \simeq [0.56, 0.88]. \tag{D.7}$$

The **central value** of the ratio is $\simeq 0.72 \pm 0.16$. The upper value 0.88 is used in §7.1 as the "tight concatenation" prediction; a measurement of the ratio at or below 0.5 would be consistent with the lower band.

### *D.6 Uncertainty estimate*

The primary uncertainties are: - CP-projector factor (0.3–0.7 range); - RG running from $\mu_B = 74$ GeV to $m_b = 4.18$ GeV for $\mathcal{O}_{10}$ (small, $\lesssim 5\%$); - Higher-order non-perturbative corrections to the $\theta_{EW}$ matching (harder to estimate, possibly $\pm 20\%$).

A rigorous derivation beyond leading order is deferred to a future Paper ε on topological flavor contributions.

---

## Appendix E — Red Team Self-Audit (internal, pre-Vega external)

**R1.** *Is the geometric budget $\leq 0.55$ clearly stated as upper bound, not central prediction?* **Response.** Yes; §7.4 and §10 explicitly emphasize "upper bound, not central prediction". The abstract uses "bounded above" language.

**R2.** *Are the four falsification criteria clearly tabulated?* **Response.** Yes; §8 provides the unified Pre-Registration Card (PRC-FCNC-v2.1) table.

**R3.** *Is Channel E' still labeled as hypothesis, not theorem?* **Response.** Yes; §5.1 explicitly states "at present this channel is a structured geometric hypothesis, not yet a theorem."

**R4.** *Is Channel N structural closure rigorous?* **Response.** Yes, via Paper β Theorem 5.1. Referenced in §6.

**R5.** *Is $\Delta C_{10}^{(K)}$ derivation now explicit?* **Response.** Yes; Appendix D gives the full matching calculation with uncertainty range.

**R6.** *Is channel accounting uniform?* **Response.** Yes; §3.3 gives precise counts (10 E/N, 1 SC, 4 PV, 5 NA; total 20 entries including Q, R as non-applicable framework extensions).

**R7.** *Is the LHCb apparent-value language sufficiently prudent?* **Response.** Yes; §1.1 and throughout, the terms "apparent central value", "subject to hadronic reinterpretation", "simplified one-phase interpretation" are used consistently.

**R8.** *Anti-numerology gate.* All inputs derived or PDG. Budget is a sum of independent bounds, not a fit. Appendix D shows the Channel K calculation explicitly.

---

## References


1. Calzighetti, S. & Lucy (2026). *Paper α: Chiral Vacuum Selection and Anti-S-Duality in the 3D+3D Framework*, v1.3.
2. Calzighetti, S. & Lucy (2026). *Paper β: Closure of Anti-S-Duality in the 3D+3D Framework*, v1.1.
3. Calzighetti, S. & Lucy (2026). *Paper XCVI: The Bridge Theorem*, v1.4.
4. Calzighetti, S. & Lucy (2026). *Paper LIV: Three Generations from 6D Orbifold*, v1.3 (FINAL).
5. Calzighetti, S. & Lucy (2026). *Paper FCNC at the Bridge Scale*, v1.0 and v2.0.
6. Alguero, Bordone, Biswas, Crivellin et al. (2025). *Global fit to $b \to s\ell\ell$ transitions after LHCb Run 1+2 and CMS 2025*. Forthcoming.
7. Inami, T. & Lim, C. S. (1981). Prog. Theor. Phys. 65, 297.
8. Buchalla, G., Buras, A. J., Lautenbacher, M. E. (1996). Rev. Mod. Phys. 68, 1125.

9. Khodjamirian, A., Mannel, T., Pivovarov, A. A., Wang, Y.-M. (2010). JHEP 09, 089.
10. Ciuchini, M., Fedele, M., Franco,