The Single-Scheme χ Audit

Testing the Up/Down Susceptibility Profile under Consistent MS̄ Running — and Why ρ = 0.50 ± 0.02 Matters

Keith Taylor *VERSF Theoretical Physics Programme — Quark Mass Hierarchy Series*

Audit Verdict

G-CHI-1 removes the QCD/MS̄ scheme-mixing worry. Under consistent MS̄ QCD running with four-loop threshold matching, the χ-profile ratio is

$\rho = 0.503 \pm 0.023$ (Monte Carlo: median 0.502, 68% interval [0.475, 0.526], $N = 10^5$),

fully consistent with ½ — the value 0.500 sits at the 47th percentile of the input-uncertainty distribution ($P(\rho \geq 0.500) = 53\%$), i.e. essentially central. The audit is **consistent with the structurally-predicted ½ and fails to refute it**; it does **not** measure exact equality, and — because the ±0.02 band admits a continuum of nearby ρ — it does not discriminate ½ from its neighbours. The 0.503-vs-0.500 distinction lies below the present input precision (dominated by m_u/m_d; note the headline 0.503 uses the *individual-mass* ratio m_u/m_d = 0.4596 — the better-constrained direct lattice ratio 0.474 gives 0.508, so the central value is choice-dependent over ≈ [0.503, 0.508], both well inside the band). **Reproducibility:** the headline is re-run reproducible (same wrapper, same inputs); an independent-implementation check is still outstanding — the one genuinely independent attempt on record, the 0.54 branch, disagreed and is explained but not yet diagnosed (§11). **Scope:** this is a QCD/MS̄ audit; QED/electroweak running is not included — a back-of-envelope estimate (Appendix C) puts its effect on ρ at $\delta\rho_{QED} \approx +0.002$ (band 0 to ≈ +0.0025), subdominant to the input band but comparable to the 0.003 excess, hence required before any ρ ≈ 0.001 test. The structural origin of ½ — the Closure-Order Twin and Log-Access Intertwining Lemmas — remains to be derived.

Summary for the General Reader

The previous paper studied a striking pattern in the quark masses. In each generation there is an up-type quark and a down-type quark, and the quantity χ measures the log-ratio between them:

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

Using the standing values from the quark mass-ratio grid, the *increase* in χ from generation 1 to generation 2 looked about twice the increase from generation 2 to generation 3 — the increments appeared to halve, $\rho = \Delta\chi_2/\Delta\chi_1 \approx ½$. That halving matters because the programme has a structural candidate for it: a two-fold up/down structure in the G sector, which the previous paper argued would force exactly ½ *if* the two folds are a matched "twin" pair feeding a fixed averaging operation.

But that paper flagged one serious danger before any of the structure could be trusted. The numbers being compared were not all quoted in one clean mass convention — some at 2 GeV, some at the top scale, and the top/bottom ratio could accidentally mix two different definitions of quark mass ("pole" and "$\overline{\text{MS}}$"). Across those choices the apparent ρ could swing anywhere from about 0.39 to 0.50, so the near-½ pattern might have been a bookkeeping artefact rather than physics. The previous paper therefore named one job as the highest priority: recompute all three χ points in one consistent convention at one common scale, and see whether ½ survives.

This paper does that job, using the standard QCD tool RunDec to bring all six quark masses to a common scale in a single consistent ($\overline{\text{MS}}$) treatment, with four-loop running and proper threshold matching. The headline:

$\rho(M_Z) \approx 0.5030$, $\rho(m_t) \approx 0.5033$.

Three things follow, and the order matters. **First, the QCD/$\overline{\text{MS}}$ scheme-mixing worry is resolved.** The clean audit reproduces the near-½ value rather than destroying it; the old number was not a mixed-convention accident. **Second, the result is remarkably stable** — moving the common scale from M_Z to the top mass changes ρ by only 0.0003, because consistent running cancels almost completely in these ratios. That stability is the genuinely solid finding. **Third — and this is the honest correction to any triumphant reading — the central value 0.503 is *not* meaningfully "0.7% above ½," because the uncertainty on ρ is far larger than 0.7%.** The dominant inputs, the up and down quark masses, are among the worst-measured numbers in particle physics, and they alone put an uncertainty of roughly ±0.02 on ρ. So the proper statement is: $\rho = 0.50 \pm 0.02$, fully consistent with ½, with the small central excess well inside the noise.

This has a sharp consequence for the theory. The structural mechanism predicts *exactly* ½, with a known way to fail: if the two folds are slightly "adjacent" rather than perfect twins, ρ drops below ½ by a calculable amount. The current data cannot test this, because the precision needed to distinguish "exactly ½" from "½ minus a small adjacency" is about twenty times finer than what the up/down mass uncertainty currently allows. The audit has therefore done exactly what it was meant to do — it has removed the scheme worry and confirmed the target is the right one — while making clear that turning ρ into a real test of the mechanism now waits on better light-quark masses, not on more theory.

To place this result in the whole picture, here is the full quark grid the programme works from — every quark mass relative to the down quark — with the VERSF structural estimate, the measured value, and the masses *this* audit computes under one consistent $\overline{\text{MS}}$ convention:

relative to down = 1	**down**	**up**	**strange**	**charm**	**bottom**	**top**
VERSF structural estimate	1	0.46	20	235	1,040	61,500
measured (rounded)	1	0.47	20	235	1,080	64,000
this audit — clean $\overline{\text{MS}}$ at m_t	1	0.46	≈ 20	≈ 232	≈ 1,060	≈ 63,000
how close (estimate vs measured)	—	consistent	imported	< 1%	4%	4%
status	anchor	conditional prediction	import (resolution limit)	match awaiting χ law	prediction	conditional by-product

Two things to read from it. The audit row — the masses this paper actually ran, expressed relative to down — falls between the structural estimate and the measured value in every column, so the consistent $\overline{\text{MS}}$ run reproduces the whole grid and confirms (as it did for ρ) that the grid's measured values are not a scale-mixing artefact. And the *status* row is **unchanged** by this paper: down is the anchor; up is a prediction conditional on the K = 7 form; strange is imported at the resolution limit; charm and top are matches awaiting the structural χ law; bottom is a genuine channel-saturation prediction. This audit sharpens the numbers, not the status labels — its own deliverable is just the three same-generation up/down ratios (up, charm/strange, top/bottom) whose increment ratio is $\rho = 0.503 \pm 0.023$.

The honest status: **the empirical target survives a clean audit; ½ is confirmed as consistent; the precision is now limited by the up/down mass inputs, not by convention choices; and the structural theorem remains to be proved.**

Table of Contents

- Summary for the General Reader
- Abstract
- 1. Motivation and the Question This Paper Settles
- 2. Definitions
- 3. Audit Standard
- 4. Inputs
- 5. Computational Method
- 6. Validation, Loop-Order Convergence, and Reproducibility
- 7. Result at $\mu = m_t$
- 8. Result at $\mu = M_Z$
- 9. Scale and Scheme Stability — the Solid Finding
- 10. The Uncertainty Budget — Why 0.503-vs-0.500 Is Not Yet a Measurement
- 11. The Charm/Strange Systematic and the 0.54 Branch
- 12. Comparison with the Old Mixed-Scheme Value
- 13. Structural Reading — What $\rho \approx 0.503$ Says About the Twin Mechanism
- 14. Impact on the χ-Profile Theorem
- 15. What This Paper Does and Does Not Establish
- 16. Status Ledger
- 17. Conclusion
- Appendix A — Reproducibility Skeleton
- Appendix B — Monte Carlo Error Propagation
- Appendix C — QED/Electroweak Differential: a Back-of-Envelope

---

Abstract

The preceding paper, *Binary Fold Resolution and the χ Column*, isolated the same-generation up/down quark mass split $\chi(g) = \ln(m_{up}(g)/m_{down}(g))$ as the live target of the VERSF quark hierarchy programme, with increments $\Delta\chi_1 = \chi(2) - \chi(1)$, $\Delta\chi_2 = \chi(3) - \chi(2)$ and profile ratio $\rho = \Delta\chi_2/\Delta\chi_1$ apparently ≈ ½. It warned that the target value itself was insecure, because the quoted χ points mixed scales and potentially mixed mass schemes (pole vs M$\bar{\text{S}}$), across which ρ ranges ≈ 0.39–0.50. It named the single-scheme, single-scale audit **G-CHI-1** as the highest-priority next action.

This paper performs that audit with RunDec/CRunDec, four-loop QCD running and M$\bar{\text{S}}$ threshold matching, from PDG-style M$\bar{\text{S}}$ inputs, bringing all six masses to common scales M_Z and m_t. The result (re-run reproducible; an independent-implementation cross-check is still outstanding, §6/§11):

$\rho(M_Z) \approx 0.5030$, $\rho(m_t) \approx 0.5033$.

The profile remains close to the exact-halving target under consistent running. Three findings frame the result. **(1) Scheme/scale stability.** Under *consistent* $\overline{MS}$, ρ is nearly scale-independent — the M_Z↔m_t drift is 0.0003 — because the QCD mass anomalous dimension is flavour-universal and common running cancels in same-generation ratios; the 0.39 end of the old band came specifically from *mixing* pole and $\overline{MS}$, which the audit removes. **(2) An uncertainty budget, previously absent.** $\rho = 0.503 \pm 0.023$ (input-dominated), with ≈ 83% of the variance from m_u/m_d alone and the next contributions from the charm/strange pair; truncation (3→4 loop) contributes ≈ 0.002. The central excess of 0.003 above ½ is ≈ 7× smaller than the input uncertainty and is *not* physically resolvable. **(3) A structural reading.** The twin mechanism predicts ρ = ½ exactly, with a fold-flip edge of weight β giving $\lambda_{odd} = ½(1-\beta)$; the observed ρ corresponds to $\beta = 1 - 2\rho \approx -0.007$ (≈ −0.016 with the direct lattice ratio). The central position above ½ is real and robust to loop order (the series converges from below to $\rho_\infty \approx 0.5037$, §6) — but its interpretation is genuinely multi-way ambiguous (physical β < 0; perturbative structure; the m_u/m_d central choice; an uncomputed QED piece of ≈ +0.002, §9/Appendix C), each comparable to the 0.003 excess. The defensible statement is therefore only that β = 0 is consistent within ±0.05 and the sign of the central value carries no information. Resolving β at the 0.007 level requires $\sigma_\rho \approx 0.001$, ≈ 20× the present precision, bounded by m_u/m_d.

Status: [G-CHI-1 — QCD/$\overline{MS}$ scheme-mixing worry resolved; target ½ confirmed consistent under clean $\overline{MS}$ running.] The central value is ρ = 0.503 ± 0.02, consistent with but not measurably distinct from ½; the 0.503-vs-0.500 distinction is input-limited, not derived. The structural theorem remains conditional on the Closure-Order Twin Lemma and the Log-Access Intertwining Lemma. The audit confirms the programme is aiming at the right empirical shape; it does not prove the mechanism.

---

1. Motivation and the Question This Paper Settles

The VERSF quark mass-ratio programme separates the hierarchy into a same-charge baseline ladder B(g) and a same-generation up/down susceptibility χ(g) = ln(m_up(g)/m_down(g)), with

χ(1) = ln(m_u/m_d), χ(2) = ln(m_c/m_s), χ(3) = ln(m_t/m_b).

The previous paper proposed a geometric increment profile $\Delta\chi_{g+1} = \rho\cdot\Delta\chi_g$ with the structural target ρ = ½, sourced by binary fold resolution in the two-dimensional G sector. The strongest result there was explicitly *conditional*: **if** the physical fold pair is a nonadjacent twin pair under the access graph (the Closure-Order Twin Lemma), and **if** the microscopic access dynamics descend to the half-lazy operator on the additive log-access shell (the Log-Access Intertwining Lemma), then the χ increments halve exactly. The mechanism predicts ρ = ½ with no error term.

That paper also issued a prior empirical warning, and made it the top-priority gate. The standing χ values were assembled from masses quoted at different scales and — more dangerously — in potentially mixed schemes. The exact value ρ = ½ was therefore not established *even observationally*: a worked illustration showed ρ swinging across ≈ 0.39–0.50 depending on whether pole and $\overline{MS}$ conventions were mixed, with the reported 0.500 sitting at the favourable end of that band. The single most important next calculation was named G-CHI-1: recompute χ(1), χ(2), χ(3) in one consistent scheme at one common high scale, and decide whether ½ is the right target at all.

This paper answers exactly that question — and, with an uncertainty budget the previous treatment lacked, says how precisely the question can currently be answered.

---

2. Definitions

$\chi(g) = \ln(m\_up(g) / m\_down(g))$, $\chi_1 = \ln(m\_u/m\_d)$, $\chi_2 = \ln(m\_c/m\_s)$, $\chi_3 = \ln(m\_t/m\_b)$, $\Delta\chi_1 = \chi_2 - \chi_1$, $\Delta\chi_2 = \chi_3 - \chi_2$, $\rho = \Delta\chi_2/\Delta\chi_1$.

The exact-halving target is ρ = ½. The goal of this audit is **not** to derive ρ but to determine, after consistent QCD running and threshold matching, (i) whether the empirical value remains near ½, (ii) how stable it is to scale and scheme, and (iii) how precisely it is currently known.

---

3. Audit Standard

The audit meets five requirements.

1. **One convention.** All six quark masses expressed as $\overline{MS}$ running masses.
2. **Common scale.** Masses compared at a common scale, not mixed across 2 GeV, m_b, m_t without correction. Because the top mass cannot be run to 2 GeV, the common scale is a *high* scale — M_Z and m_t.
3. **Explicit thresholds.** Decoupling handled explicitly at m_c, m_b, m_t.
4. **Fixed perturbative order.** Four-loop running and matching throughout; loop-order stability reported (§6).
5. **Exposed intermediates.** All intermediate masses and ratios printed, because the one materially sensitive ratio is c/s, which straddles the charm threshold.

---

4. Inputs

PDG-style $\overline{MS}$ inputs, with approximate present uncertainties (used in the budget of §10):

Quantity	Central	Approx. uncertainty
α_s(M_Z)	0.1180	±0.0009
m_u(2 GeV)	2.16 MeV	$^{+0.49}_{-0.26}$ MeV

Quantity	Central	Approx. uncertainty
m_d(2 GeV)	4.70 MeV	$^{+0.48}_{-0.17}$ MeV
m_s(2 GeV)	93.5 MeV	≈ ±1–2%
m_c(m_c)	1.27 GeV	±0.02 GeV
m_b(m_b)	4.18 GeV	±0.03 GeV
m_t(m_t)	162.5 GeV	≈ ±0.5 GeV
M_Z	91.1876 GeV	—

Thresholds taken at the corresponding $\overline{\text{MS}}$ masses (m_c = 1.27, m_b = 4.18, m_t = 162.5 GeV). The combination m_u/m_d is better constrained by lattice as a *direct ratio*, m_u/m_d = 0.474 $^{+0.056}_{-0.074}$; the audit's central inputs give m_u/m_d = 0.4596, ≈ 3% below that ratio central — a difference tracked in §10.

---

5. Computational Method

The audit uses `rundec v0.7` (Python binding to CRunDec/RunDec). The strong coupling is run with `AlphasExact` and decoupled at flavour thresholds with `DecAsDownMS`/`DecAsUpMS`; masses are run with `mMS2mMS` and matched with `DecMqDownMS`/`DecMqUpMS`. Two common-scale forms are computed:

1. μ = m_t, all masses brought to the six-flavour convention at the top scale;
2. μ = M_Z, the natural five-flavour convention for u, d, s, c, b with a six-flavour run for the top, plus a six-flavour cross-check.

Because QCD running is flavour-universal *within* a fixed-flavour theory, any ratio of two masses both active in that theory (e.g. u/d, c/s within the same n_f) is essentially scale-invariant. The only materially sensitive ratio is c/s, because charm and strange sit on opposite sides of the charm threshold — so all sensitivity to matching choices concentrates there. This is not a nuisance to be hidden; it is the single lever that the budget (§10) and the systematic check (§11) both turn on.

---

6. Validation, Loop-Order Convergence, and Reproducibility

Three checks precede the result.

Standard validation. At μ = 4.18 GeV with n_f = 5,

$\alpha_s(4.18\ \text{GeV}) = 0.224616$, consistent with the expected ≈ 0.225. Running $m_b(m_b) = 4.18$ GeV upward,

$m_b(M_Z) = 2.863070$ GeV,

within the expected $\overline{MS}$ range ≈ 2.75–2.86 GeV.

Loop-order convergence. Recomputing $\rho(m_t)$ at successive orders:

2-loop: $\rho = 0.49316$, 3-loop: $\rho = 0.50165$, 4-loop: $\rho = 0.50331$.

The series is **monotone increasing and converges from below**, with shifts +0.00849 then +0.00166 (ratio ≈ 0.20). A geometric extrapolation of the tail puts the all-orders value at $\rho_\infty \approx 0.5037$ — *above* the four-loop number, i.e. slightly **further** from ½, not closer. Two consequences, the second correcting an easy error. The residual truncation *uncertainty* is small — the geometric tail beyond four loops is only $\rho_\infty - \rho_4 \approx 0.0004$ — but §10 conservatively carries the larger last-shift value, ±0.002 (the full 3→4-loop step), as the budget's truncation bound rather than the tail estimate. Either way truncation does **not** explain the central excess above ½: going to all orders widens the gap, so the ≈ 0.003 excess is a robust perturbative feature, not an artefact that vanishes at higher order. What makes ½ consistent is the ±0.023 input band of §10, not truncation.

Reproducibility — and the limit of it. The full four-loop pipeline (α_s chain, light/charm/bottom running and matching to m_t) was re-run from the inputs of §4 and reproduces the headline to the last quoted digit: $\alpha_s(4.18) = 0.224616$, $m_b(M_Z) = 2.863070$, and $\rho(m_t) = 0.50331$ with the masses and ratios of §7. This confirms the result is not a transcription or single-run artefact — but it is a *re-run through the same RunDec/CRunDec wrapper from the same inputs*, so it tests transcription and run-error, **not** tool or methodology independence. A genuinely independent check (a separate CRunDec implementation reproducing c/s ≈ 11.68, per the §11 falsification standard) remains outstanding; the one independent attempt on record, the 0.54 branch, *disagreed* and is explained as a c/s threshold systematic (§11) but has not been diagnosed line-by-line. The headline's reproducibility and that undiagnosed disagreement are stated together, here, so the claim is not read as stronger than it is.

---

7. Result at $\mu = m_t$

Running all six masses to $\mu = m_t$:

Quantity	Value at m_t
m_u	0.001186 GeV
m_d	0.002581 GeV
m_s	0.051341 GeV
m_c	0.599728 GeV

Quantity	Value at m_t
m_b	2.729894 GeV
m_t	162.500000 GeV

Ratio	Value
m_u/m_d	0.459574
m_c/m_s	11.681278
m_t/m_b	59.526113

$\chi_1 = -0.777454$, $\chi_2 = 2.457987$, $\chi_3 = 4.086415$, $\Delta\chi_1 = 3.235442$, $\Delta\chi_2 = 1.628428$, **ρ(m_t) = 0.50331.**

Close to ½, not exactly ½.

---

8. Result at μ = M_Z

At μ = M_Z (five-flavour convention for u, d, s, c, b; top run down in the six-flavour theory):

Quantity	Value at M_Z
m_u	0.001244 GeV
m_d	0.002707 GeV
m_s	0.053846 GeV
m_c	0.628985 GeV
m_b	2.863070 GeV
m_t	170.267433 GeV

Ratio	Value
m_u/m_d	0.459574
m_c/m_s	11.681278
m_t/m_b	59.470224

$\chi_1 = -0.777454$, $\chi_2 = 2.457987$, $\chi_3 = 4.085476$, $\Delta\chi_1 = 3.235442$, $\Delta\chi_2 = 1.627488$, **ρ(M_Z) = 0.50302.**

A six-flavour cross-check (all masses transported to n_f = 6 and run to M_Z) returns ρ ≈ 0.50331, since common QCD running cancels in the ratios.

---

9. Scale and Scheme Stability — the Solid Finding

Note that u/d and c/s are *identical* at both scales (0.459574 and 11.681278 to six figures), and only t/b moves — from 59.526 at m_t to 59.470 at M_Z, a 0.09% shift. This is not luck; it is the structure of the problem. The QCD mass anomalous dimension $\gamma_m(\alpha_s)$ is flavour-universal, so within a fixed flavour theory a ratio of two active masses does not run at all. Every scale- and scheme-dependence of ρ is therefore confined to *threshold matching* — the points where one member of a ratio is "heavy" and the other "light" across a flavour boundary (s vs c across the charm threshold; b vs t across the top threshold).

The consequence is the genuinely robust result of this audit:

$\rho(M_Z) = 0.50302$, $\rho(m_t) = 0.50331$, $|\Delta\rho| = 0.00029$.

Under *consistent* $\overline{MS}$, the scale-dependence of ρ is $\approx 3 \times 10^{-4}$ — two orders of magnitude smaller than the 0.39–0.50 band the previous paper feared. That band was real, but it was a **scheme-mixing** artefact: it arose specifically from comparing a pole-mass top against an $\overline{MS}$ bottom. Once pole and $\overline{MS}$ are not mixed, the scale-and-scheme freedom in ρ collapses to the percent-of-a-percent level. **The "consistency, not precision" hedge of the previous paper is, for the scale/scheme axis, now discharged: ρ is scheme-stable.** What remains uncertain is not the convention but the inputs — which is the subject of §10.

[Established here — ρ is scale/scheme-stable under consistent $\overline{MS}$; the old 0.39–0.50 band was pole-vs-$\overline{MS}$ mixing, now removed.]

Scope — this is a QCD/$\overline{MS}$ audit (the QED/electroweak caveat, now estimated). RunDec/CRunDec is a QCD running and decoupling tool, and that is exactly what the dominant ambiguity demanded: the audit removes the pole-vs-$\overline{MS}$ scheme mixing and the QCD threshold-matching bookkeeping that drove the old 0.39–0.50 band. It does **not** include QED/electroweak running. This is the one effect flavour-universal QCD cancellation does *not* protect, because the squared charges differ: $Q^2_{up} = 4/9 \neq 1/9 = Q^2_{down}$, so the differential QED mass anomalous dimension is nonzero, $\approx (3/2)(Q^2_{up} - Q^2_{down})(\alpha/\pi) = (1/2)(\alpha/\pi)$ per unit $\ln \mu$. Rather than assert this is negligible, Appendix C estimates it. The per-generation differential shifts on χ are $\delta\chi_{QED}(g) \approx -0.006$ to $+0.001$, and they would cancel almost exactly in ρ — *if* the three up/down pairs shared a common reference scale. They do not: the audit references the top at m_t (zero QED lever) while the bottom runs up from m_b, which breaks the cancellation and leaves a net **$\delta\rho_{QED} \approx +0.002$** (computed +0.0023; band ≈ 0 for an idealized common reference scale to $\approx +0.0025$ across mass-definition and α choices, positive in the audit's scheme). So the honest statement is not "negligible" but bounded: QED is **subdominant to the ±0.023 input band** (the headline is safe), yet it is **comparable to the truncation shift and to the 0.003 excess above ½, and carries the same sign** — so it is *not* negligible at the $\rho \approx 0.001$ β-test frontier of §13, where it must be included before any claim about β. The exact size and sign depend on the low-scale mass-definition choices, which is itself the reason a full treatment is the correct next refinement.

[Scope note — QCD/$\overline{MS}$ audit; QED/EW not included. Estimated effect on ρ is $\delta\rho_{QED} \approx +0.002$ (band 0 to $\approx +0.0025$, Appendix C): subdominant to the input band but comparable to the excess above ½ and to truncation, hence required at the $\rho \approx 0.001$ frontier, not before.]

10. The Uncertainty Budget — Why 0.503-vs-0.500 Is Not Yet a Measurement

The previous draft reported $\rho \approx 0.503$ and read the 0.003 gap above ½ as "0.7% high." That reading is not yet warranted, because no uncertainty was attached. ρ is a ratio of differences of logs of masses, and it inherits the (substantial) uncertainties of the worst-measured masses in the Standard Model. This section supplies the budget the conclusion requires.

Sensitivities. With $\rho = (\chi_3 - \chi_2)/(\chi_2 - \chi_1)$,

$\partial\rho/\partial\chi_1 = +\rho/\Delta\chi_1 = +0.156$, $\partial\rho/\partial\chi_2 = -(1+\rho)/\Delta\chi_1 = -0.465$, $\partial\rho/\partial\chi_3 = +1/\Delta\chi_1 = +0.309$.

χ_2 carries the largest coefficient, but χ_1 carries the largest *uncertainty*, and the second wins.

Per-input contributions (propagated through the full RunDec pipeline where the input enters running/matching, otherwise analytically):

Source	Input uncertainty	δρ
m_u/m_d	σ(ln u/d) ≈ 0.137 (direct-ratio lattice)	±0.021
m_c(m_c)	±0.02 GeV (±1.6%)	±0.007
m_s(2 GeV)	±1.3%	±0.006
m_b, m_t	±0.7%, ±0.3%	±0.002
truncation	3→4 loop	±0.002
scale/scheme	M_Z ↔ m_t	±0.0003

Added in quadrature (treating the inputs as uncorrelated — a conservative choice; see below):

σ_ρ ≈ 0.023, dominated ≈ 83% by m_u/m_d,

so

ρ = 0.503 ± 0.023 (input) ⟹ ρ = 0.50 ± 0.02.

Two consequences. **(1) ½ is consistent at well under 1σ.** The central excess of 0.003 is ≈ 7× smaller than σ_ρ; on the present inputs it is not distinguishable from zero. (It is *not* explained by truncation: §6 shows the loop series converges from below to ρ_∞ ≈ 0.5037, so higher orders widen the gap rather than close it — the excess is consistent with ½ only through the input band, not through perturbative incompleteness.) **(2) The light-quark masses are the binding constraint.** Better m_c, m_s, m_t/m_b would barely move σ_ρ; only a sharper m_u/m_d does, because it carries 83% of the variance.

On correlations and central choices. The budget above treats inputs as independent, which over-counts: lattice determines m_u/m_d as a direct ratio better than the individual masses, so the true σ(ln u/d) may be somewhat below 0.137, tightening σ_ρ modestly — but m_u/m_d remains dominant either way. The *central* choice matters too, and the headline rests on the **less well-constrained** of the two determinations: this paper quotes ρ = 0.503 from the individual-mass ratio m_u/m_d = 0.4596 (for continuity with the standing grid), but the lattice *direct ratio* m_u/m_d = 0.474 is the better-measured input, and it gives ρ = 0.508. So the central value is choice-dependent over ≈ [0.503, 0.508]; by the precision argument the natural headline is arguably 0.508 with 0.503 as the variant, not the reverse. Both lie well inside the band, and the move (≈ 0.005) is what the §13 reading must keep in view. A publication-grade number should carry the full correlated lattice covariance; the conclusion (ρ = 0.50 ± 0.02, ½ consistent, input-limited) is robust to that refinement.

Cross-check by Monte Carlo. The linearized budget is confirmed by a 10^5-sample Monte Carlo over all inputs with asymmetric light-quark errors and matching-scale jitter (Appendix B): median ρ = 0.502, 68% interval [0.475, 0.526] (σ ≈ 0.026), with 0.500 at the 47th percentile (P(ρ ≥ 0.500) = 53%) — essentially central, marginally below the median. The MC σ is slightly larger than the linearized 0.023 — the difference is the asymmetric m_u/m_d tail and the matching jitter — and the MC median sits marginally below the central point value for the same asymmetry reason. Both methods agree on the operative conclusion.

[Established here — ρ = 0.50 ± 0.02 (linearized and Monte Carlo agree), input-dominated (83% from m_u/m_d); ½ consistent at < 1σ, with 0.500 essentially at the MC median (47th percentile); the 0.503-vs-0.500 gap is not resolvable at current precision, and the central value is choice-dependent over [0.503, 0.508].]

---

11. The Charm/Strange Systematic and the 0.54 Branch

An independent attempt at this calculation initially produced ρ ≈ 0.540 (M_Z) / 0.542 (m_t). The present consistent-stepwise audit does not reproduce those values, and the reason is diagnostic rather than mysterious: **it is most naturally explained as the size of the charm/strange threshold systematic, not a separate physical result.** To reach ρ ≈ 0.54 while keeping u/d and t/b near their audited values, the middle ratio c/s must fall from ≈ 11.68 to ≈ 10.8; that shifts χ_2 by −0.078 and, via $\partial\rho/\partial\chi_2$ = −0.465, moves ρ by +0.037 — exactly the 0.503 → 0.54 gap.

Matching-scale variation. To show that the audited value is stable under *reasonable* charm-matching choices — and that the outliers are identifiable edge cases — the charm decoupling scale μ_match,c was varied directly (all else held at central inputs):

Charm matching scale μ_match,c	m_c/m_s (at m_t)	ρ
m_c / 2 (= 0.635 GeV, *below* the charm mass; α_s large, 4-loop matching unreliable)	12.90	0.4585

Charm matching scale μ_match,c	m_c/m_s (at m_t)	ρ
m_c (central)	**11.68**	**0.5033**
2 m_c	11.71	0.5023
diagnostic "bad branch" (mismatched c/s)	≈ 10.8	≈ 0.541

Over the physically sensible band μ_match,c ∈ [m_c, 2m_c], ρ is stable to ≈ 0.001 (0.5023–0.5033). Leaving that band requires either matching *below* the charm mass (μ = m_c/2, where the perturbative series is no longer trustworthy and ρ drops to 0.46) or an outright threshold mishandling (the bad branch, ρ ≈ 0.54). So **0.503 is stable under reasonable choices; both 0.54 and 0.46 are identifiable edge/erroneous cases, not competing results.** The matching-scale jitter over [m_c, 2m_c] is folded into the Monte Carlo of Appendix B and contributes negligibly.

The candidate mishandlings that inflate ρ toward 0.54, all in the charm/strange sector:

1. running m_c or m_s with α_s from the wrong flavour theory;
2. failing to match m_s upward correctly through the charm threshold;
3. treating m_c(m_c) as belonging to the wrong n_f theory;
4. applying a decoupling function in the wrong direction;
5. using a fixed-n_f shortcut for one member of the c/s pair but not the other.

The one-line diagnostic: print m_c(μ), m_s(μ), and c/s before computing χ. **c/s ≈ 11.68 ⟹ ρ ≈ 0.503; c/s ≈ 10.8 ⟹ ρ ≈ 0.54; c/s ≈ 12.9 ⟹ ρ ≈ 0.46.** The consistent stepwise treatment of §5–§8 gives 11.68 and is the one whose scale-stability (§9) is 3×10^{-4}.

What would falsify this audit. The result would be overturned if an independent RunDec/CRunDec implementation — same input masses, loop order, threshold scales, and $\overline{\text{MS}}$ conventions — produced c/s far from 11.68 at m_t, or ρ outside the quoted band. The decisive diagnostic is the single sensitive ratio: agreement on c/s is necessary and nearly sufficient for agreement on ρ, since u/d and t/b are scale-stable and weakly levered. Any reproduction attempt should therefore **print c/s before ρ**. The audit is falsifiable at exactly this point, and that is the right standard — note that the §6 re-run is *not* such an independent check (same wrapper, same inputs), and the one genuinely independent attempt on record (the 0.54 branch) returned c/s ≈ 10.8 and remains explained-but-not-diagnosed. So the independent-implementation check is the outstanding test, and the falsifier is a concrete number (c/s ≈ 11.68 at m_t), not a matter of interpretation.

[Not reproduced; most naturally explained by charm/strange threshold handling. The 0.54 branch should not presently be treated as a competing physical value; it is best understood as a charm/strange threshold-handling branch (c/s lever ≈ 0.037) pending line-by-line independent diagnosis — the exact mishandling that produced it has not been isolated. Consistent stepwise matching gives ρ ≈ 0.503, stable to ≈ 0.001 over μ_match,c ∈ [m_c, 2m_c].]

12. Comparison with the Old Mixed-Scheme Value

Ratio	Old grid (mixed)	This audit (clean $\overline{MS}$)
m_u/m_d	0.4615 (= 6/13)	0.4596
m_c/m_s	11.76	11.681
m_t/m_b	59.4	59.5
ρ	**≈ 0.500**	**≈ 0.503**

The consistent audit does not destroy the near-halving profile; it reproduces it, shifting ρ by +0.003 — well within the §10 budget. The old mixed-scheme value was not wildly misleading. The clean statement is **not** "ρ = ½ proven" but "ρ is close to ½ under consistent $\overline{MS}$ running, and the closeness is not a convention accident."

13. Structural Reading — What ρ ≈ 0.503 Says About the Twin Mechanism

This is the point at which the audit feeds back into the structural programme, and it is sharper than the previous paper noticed it could be.

The twin-contraction theorem (the previous paper's §6M.1.3) predicts ρ = ½ **exactly** — there is no error term in $\ell_{odd} \circ W = \frac{1}{2}\cdot\ell_{odd}$. The same paper also derived the *failure mode*: a direct fold-flip edge of normalised weight β between the two folds gives

$\lambda_{odd} = \frac{1}{2}(1 - \beta)$, i.e. $\rho = \frac{1}{2}(1 - \beta)$.

So the measured ρ maps directly onto β:

$\beta = 1 - 2\rho$.

At the audited central value ρ = 0.5033, this is **β ≈ −0.007** (and β ≈ −0.016 at ρ = 0.508, the direct-ratio central). Two observations, neither of which the previous paper could make before G-CHI-1 returned a clean number — but the first must be read with no directional bias:

(1) The central value sits above ½, and that is genuinely uninterpretable. A fold-flip edge has β > 0 and lowers ρ below ½; the audited ρ sits slightly *above* ½ (β < 0). It is tempting to read this as disfavouring adjacency — but that reading does not survive scrutiny, and the paper does not make it. The excess above ½ is consistent with at least four sources, each comparable to or larger than the 0.003 excess itself: (a) a genuine physical β < 0; (b) **robust perturbative structure** — the loop series converges from below to $\rho_\infty \approx 0.5037$ (§6), so the excess is *not* a fluctuation toward exact ½ and in fact grows at higher order; (c) the **input-central choice** —

adopting the better-measured direct lattice ratio moves ρ to 0.508, $\beta \approx -0.016$, i.e. *more* negative (§10); and (d) **omitted QED**, a charge-dependent differential that does not cancel between up- and down-type and is estimated at $\delta\rho_{QED} \approx +0.002$ of the *same* sign (§9, Appendix C) — so that subtracting the envelope estimate would move the central value *toward* ½ (0.503 → ≈ 0.501), tentatively and within the crudeness of the estimate, not as a correction to be banked. With the central value pulled in the same direction by at least three uncomputed or convention-dependent effects of the same size, the only defensible statement is: **$\beta = 0$ is consistent within ±0.05, and the sign of the central value carries no information.** The exact-twin ($\beta = 0$) hypothesis is consistent with the audit; so is weak adjacency; the data adjudicate neither.

(2) The test is firmly uncertainty-limited. To distinguish exact twins ($\beta = 0$, $\rho = 0.500$) from weak adjacency at, say, $\beta = 0.007$ ($\rho = 0.4965$) at 3σ would require

$$\sigma_\rho \lesssim 0.001,$$

roughly **20× better than the present 0.023**. Since σ_ρ is 83% m_u/m_d, this is, concretely, a demand for m_u/m_d at the ≈ 1% level — a lattice target, not a theory task. And at that precision the ≈ 0.002 QED differential (§9) is no longer ignorable, so a genuine β-test needs *both* a sharper m_u/m_d and a QED treatment. Until then, G-CHI-1 establishes that the mechanism is aiming at the right value and is consistent with the exact-twin prediction, but it **cannot** adjudicate β, and no claim about small adjacency — in either direction — should be read from the central value.

This converts ρ from a number to be admired into a quantitative, if currently blunt, instrument on the mechanism: G-CHI-1 has shown the instrument reads ½ within a needle whose width is set by the up/down mass. The previous paper built the test of its own mechanism; this paper reads it, and finds the reading consistent with exact ½ and presently far too coarse — and too contaminated by perturbative, input-choice, and QED effects of the same size — to see β at all.

[Structural reading — $\rho_{obs} \Longrightarrow \beta = 1 - 2\rho \approx -0.007$ (−0.016 at the direct-ratio central); $\beta = 0$ consistent within ±0.05; the sign carries no information (the excess is matched in size by perturbative structure, the input-central choice, and an uncomputed +0.002 QED piece, all of the same sign). Resolving β needs $\sigma_\rho \approx 0.001$ (m_u/m_d to ≈ 1%) *and* a QED treatment — ≈ 20× current precision.]

---

14. Impact on the χ-Profile Theorem

The previous paper reduced the structural claim to a conditional theorem: *if* the VERSF fold-access construction satisfies its six-part antecedent — the fold projectors (OP0), the up/down identification (OP0.5), the Closure-Order Twin Lemma, the Log-Access Intertwining Lemma / geometric-mean access law, rate-one, and the 6/13 anchor — *then* $\chi_1 = \ln(6/13)$ and $\Delta\chi_{g+1} = \frac{1}{2}\cdot\Delta\chi_g$. This audit does not touch that antecedent. It checks only whether the *target value* ½ is empirically reasonable under consistent inputs.

The verdict is favourable but bounded: **the target survives, and is now scheme-stable; the precision is input-limited.** The standing interpretation should become:

G-CHI-1: the QCD/MS̄ scheme-mixing worry is resolved; $\rho = 0.50 \pm 0.02$ is consistent with the structurally-predicted ½ and fails to refute it — to consistency, not to precision, and without discriminating ½ from nearby alternatives.

This is materially stronger than the previous status — under which $\rho = ½$ could have been a mixed-scheme mirage — without overreaching into "½ measured." The next theoretical burden returns intact to the two structural gates (§15).

---

15. What This Paper Does and Does Not Establish

Established here

- A re-run-reproducible four-loop RunDec MS̄ audit (same wrapper/inputs; independent-implementation check outstanding), with passing validation ($\alpha_s(4.18) = 0.224616$; $m_b(M_Z) = 2.863070$) and clean loop-order convergence ($0.493 \to 0.502 \to 0.503$).
- Under consistent MS̄ running and matching, $\rho(m_t) = 0.50331$ and $\rho(M_Z) = 0.50302$, with a six-flavour cross-check returning 0.50331.
- **Scale/scheme stability**: ρ is scale-stable to 3×10^{-4}; the old 0.39–0.50 band was pole-vs-MS̄ mixing and is removed by consistent treatment.
- **An uncertainty budget**: $\rho = 0.50 \pm 0.02$, $\approx 83\%$ of the variance from m_u/m_d; the 0.503-vs-0.500 gap is not resolvable at current precision.
- **A 10^5-sample Monte Carlo** (Appendix B) confirming the linearized budget: median $\rho = 0.502$, 68% interval [0.475, 0.526], with 0.500 at the 47th percentile ($P(\rho \geq 0.500) = 53\%$) — essentially central.
- The 0.54 branch should not presently be treated as a competing physical value — best understood as a charm/strange threshold-handling branch (c/s lever ≈ 0.037) pending line-by-line independent diagnosis; ρ is stable to ≈ 0.001 over $\mu_{match,c} \in [m_c, 2m_c]$.
- **Structural map** $\rho \Longrightarrow \beta = 1 - 2\rho \approx -0.007$: $\beta = 0$ consistent within ± 0.05, with the sign carrying no information (the excess above ½ is matched in size by perturbative structure, the input-central choice, and an uncomputed +0.002 QED piece); the β test is uncertainty-limited by $\approx 20\times$ (binding input: m_u/m_d).
- **A QED scope estimate** (Appendix C): the omitted charge-dependent differential shifts ρ by $\delta\rho_{QED} \approx +0.002$ (band 0 to $\approx +0.0025$) — subdominant to the input band but comparable to the excess, hence required at the β-test frontier.

Not established here

- The exact identity $\rho = ½$ (consistent with, not equal to, within ± 0.02).
- The VERSF structural theorem; the Closure-Order Twin Lemma; the Log-Access Intertwining Lemma; the geometric-mean access law as a physical microscopic VERSF law.

- The magnitude $\Delta\chi_1$.
- The first-generation anchor 6/13 as a fresh derivation rather than a relabelled upstream ratio.
- Any value of the fold-flip weight β (the data bound it only at the ±0.05 level).

---

16. Status Ledger

Object	Status after this audit
χ definition; need for single-scheme audit	Established (prior paper)
RunDec/CRunDec tool choice; α_s, m_b(M_Z) validation	Appropriate; passed
Loop-order convergence (2→3→4 loop)	0.49316 → 0.50165 → 0.50331; converges from below, ρ_∞ ≈ 0.5037; excess robust
Reproducibility of ρ(m_t)	Re-run reproducible to last digit (same wrapper/inputs); independent-implementation check outstanding
ρ at m_t / M_Z	0.50331 / 0.50302 (6-flavour cross-check 0.50331)
Scale/scheme stability under consistent MS̄	Established —
Uncertainty budget (linearized)	ρ = 0.50 ± 0.02; 83% from m_u/m_d; central choice-dependent over [0.503, 0.508]
Uncertainty budget (Monte Carlo, $N=10^5$)	median 0.502, 68% [0.475, 0.526]; 0.500 at 47th pctile (Appendix B)
Exact ρ = ½	Consistent within < 1σ; central 0.003 high, not bridged by truncation (all-orders ~0.5037)
Old mixed value ≈ 0.500	Reproduced (0.503) in clean audit, not destroyed
Reported ρ ≈ 0.54	Not reproduced; should not presently be treated as a competing physical value — a charm/strange threshold-handling branch (c/s lever 0.037) pending line-by-line independent diagnosis
Matching-scale stability	ρ ∈ [0.502, 0.503] over μ_match,c ∈ [m_c, 2m_c]; edge cases 0.46 (m_c/2) and 0.54
QED/electroweak running	Not included; estimated δρ_QED ≈ +0.002 (Appendix C); subdominant to inputs, comparable to the excess; needed at the β-frontier
Structural reading β = 1 − 2ρ	≈ −0.007 (−0.016 at direct ratio); β = 0 consistent ±0.05; sign carries no information
β test resolvability	Needs σ_ρ ≈ 0.001 (≈ 20×) and a QED treatment; binding input m_u/m_d to ≈ 1%
G-CHI-1	QCD/MS̄ scheme-mixing worry resolved; target ½ confirmed consistent; exact equality not measured

Object	Status after this audit
Closure-Order Twin Lemma; Log-Access Intertwining Lemma	Still open (structural)
Magnitude $\Delta\chi_1$	Open

---

17. Conclusion

The single-scheme χ audit was the empirical gate the previous paper required before the VERSF χ-profile theorem could be taken seriously, and it has been run, re-run reproducibly, and given the error budget the conclusion needs.

Under standard four-loop $\overline{MS}$ running and threshold matching, the profile ratio is $\rho \approx 0.503$ at both M_Z and m_t, stable to 3×10^{-4} across that scale change. Three things follow. The **QCD/$\overline{MS}$ scheme-mixing worry is resolved**: the 0.39–0.50 band the previous paper feared was a pole-vs-$\overline{MS}$ mixing effect, and under consistent treatment ρ does not wander — the near-½ value is real, not a bookkeeping accident. The **precision is input-limited, not convention-limited**: $\rho = 0.50 \pm 0.02$, with 83% of the uncertainty from m_u/m_d (and the central value itself choice-dependent over [0.503, 0.508], since the better-measured direct lattice ratio gives the upper end), so the 0.003 excess above ½ is ≈ 7× inside the noise and is not a measurement of any departure from ½. And the **structural reading is now quantitative but, at present precision, uninterpretable in sign**: the observed ρ maps to $\beta = 1 - 2\rho \approx -0.007$, consistent with the exact-twin prediction $\beta = 0$ within ±0.05 — but the central excess above ½ is matched in size by robust perturbative structure (all-orders $\rho_\infty \approx 0.5037$), by the m_u/m_d central choice, and by an uncomputed QED differential of ≈ +0.002 (Appendix C), all of the same sign, so the sign of the central value carries no information and no adjacency claim — in either direction — should be read from it. The β test is ≈ 20× too coarse and additionally needs a QED treatment.

So the audit strengthens the programme in a precise and bounded way. It does not prove the VERSF mechanism, derive the quark mass ratios, or fix the magnitude $\Delta\chi_1$. It is a consistency check: it **fails to refute** the structurally-predicted ½ and shows the old near-½ value was not a mixed-scheme accident — while, by its ±0.02 band, failing to discriminate ½ from any nearby alternative. The remaining empirical limitation is the light-quark masses, not anything the theory controls.

The standing of the programme should now read:

G-CHI-1 — QCD/$\overline{MS}$ scheme-mixing worry resolved; $\rho = 0.50 \pm 0.02$ is consistent with the structurally-predicted ½ and fails to refute it (it does not measure exact equality, nor discriminate ½ from neighbours). The structural origin of ½ remains to be derived.

The next theoretical burden returns, unchanged, to the two structural gates of the previous paper:

1. prove or refute the **Closure-Order Twin Lemma**, ℓ_odd $D^{-1}A = 0$ (folds are nonadjacent twins); and
2. prove or refute the **Log-Access Intertwining Lemma**, E_{g+1} S_g = W E_g (the physical readout realises the half-lazy operator on the log-access shell).

If those gates close, the χ profile has a structural explanation and the audited $\rho \approx ½$ is its confirmed empirical shape. If they fail, the earned result is smaller but real: the χ column follows a near-geometric profile whose ratio is, under a clean consistent audit, 0.50 ± 0.02 — close to one-half, but with its structural origin still open. Either way, the empirical target is no longer the weak link: after this audit, the weak link is once again the structure.

Appendix A — Reproducibility Skeleton

The calculation structure (verified against `rundec v0.7`; the headline below was re-run from these inputs — a transcription/run-error check, not an independent-implementation check, §6/§11). API signatures should be confirmed against the installed wrapper before final publication.

```
from rundec import CRunDec
import math

rd = CRunDec()
as_MZ, MZ = 0.1180, 91.1876
mc, mb, mt = 1.27, 4.18, 162.5
mu0 = 2.0
m_u_2, m_d_2, m_s_2 = 0.00216, 0.00470, 0.0935
m_c_mc, m_b_mb, m_t_mt = 1.27, 4.18, 162.5
L = 4

# alpha_s running and threshold matching (MS-bar)
as5_mb = rd.AlphasExact(as_MZ, MZ, mb, 5, L)
as4_mb = rd.DecAsDownMS(as5_mb, mb, mb, 5, L)
as4_mc = rd.AlphasExact(as4_mb, mb, mc, 4, L)
as3_mc = rd.DecAsDownMS(as4_mc, mc, mc, 4, L)
as3_2  = rd.AlphasExact(as3_mc, mc, mu0, 3, L)
as5_mt = rd.AlphasExact(as_MZ, MZ, mt, 5, L)

def light_to_mt(m2):
    m = rd.mMS2mMS(m2, as3_2, as3_mc, 3, L)
    m = rd.DecMqUpMS(m, as3_mc, mc, mc, 3, L)
    m = rd.mMS2mMS(m, as4_mc, as4_mb, 4, L)
    m = rd.DecMqUpMS(m, as4_mb, mb, mb, 4, L)
    m = rd.mMS2mMS(m, as5_mb, as5_mt, 5, L)
    m = rd.DecMqUpMS(m, as5_mt, mt, mt, 5, L)
    return m

def charm_to_mt(mcm):
    m = rd.mMS2mMS(mcm, as4_mc, as4_mb, 4, L)
    m = rd.DecMqUpMS(m, as4_mb, mb, mb, 4, L)
    m = rd.mMS2mMS(m, as5_mb, as5_mt, 5, L)
```

```
    m = rd.DecMqUpMS(m, as5_mt, mt, mt, 5, L)
    return m

def bottom_to_mt(mbm):
    m = rd.mMS2mMS(mbm, as5_mb, as5_mt, 5, L)
    m = rd.DecMqUpMS(m, as5_mt, mt, mt, 5, L)
    return m

m_u, m_d, m_s = light_to_mt(m_u_2), light_to_mt(m_d_2), light_to_mt(m_s_2)
m_c, m_b, m_t = charm_to_mt(m_c_mc), bottom_to_mt(m_b_mb), m_t_mt

r1, r2, r3 = m_u/m_d, m_c/m_s, m_t/m_b
chi1, chi2, chi3 = math.log(r1), math.log(r2), math.log(r3)
rho = (chi3 - chi2) / (chi2 - chi1)

# Always print the charm/strange diagnostic BEFORE trusting rho (see §11):
print("c/s =", r2, "  rho =", rho)
```

Reproduced central output at μ = m_t: u/d ≈ 0.459574, c/s ≈ 11.681278, t/b ≈ 59.526113; χ = (−0.777454, 2.457987, 4.086415); **ρ ≈ 0.50331**.

Appendix B — Monte Carlo Error Propagation

To complement the linearized budget of §10 — and to handle the asymmetric light-quark mass errors a first-order estimate cannot — ρ was propagated through the full four-loop RunDec pipeline over N = 100,000 input samples. The sampled inputs:

Input	Distribution
α_s(M_Z)	Normal(0.1180, 0.0009)
m_u/m_d	split-Normal(0.4596; σ_+ = 0.054, σ_- = 0.072) — asymmetric, lattice-ratio width
m_s(2 GeV)	Normal(93.5, 2.0) MeV
m_c(m_c)	Normal(1.27, 0.02) GeV
m_b(m_b)	Normal(4.18, 0.03) GeV
m_t(m_t)	Normal(162.5, 0.5) GeV
μ_match,c	log-uniform on [m_c, 2 m_c] — matching-scale jitter

Each sample runs the complete light/charm/bottom pipeline to μ = m_t and forms $\rho = (\chi_3 - \chi_2)/(\chi_2 - \chi_1)$. Results:

Quantity	Value
median ρ	0.5020
mean ± std	0.5007 ± 0.0256
68% interval (16th–84th)	[0.4753, 0.5258]
95% interval (2.5th–97.5th)	[0.4471, 0.5481]

Quantity	Value
P($\rho \geq 0.500$)	53.1%

Three readings. **(1) ½ is essentially central.** P($\rho \geq 0.500$) = 53% places 0.500 at the 47th percentile — marginally below the median (0.502), essentially at the centre of the distribution; the exact-halving target is not merely "consistent" but central. **(2) The MC confirms the linearized budget.** $\sigma_{MC} = 0.026$ against $\sigma_{lin} = 0.023$; the small excess is the asymmetric m_u/m_d tail plus the matching jitter, neither captured at first order, and the median sitting marginally below the point value 0.5033 has the same origin. **(3) The dominant input is unchanged.** Fixing m_u/m_d at its central value collapses the spread to $\sigma \approx 0.013$, so the up/down ratio carries ≈ 73% of the Monte Carlo variance (≈ 83% in the linearized budget, which omits the matching jitter) — the same conclusion as §10, now without a linearity assumption.

Variant — PDG direct-ratio centring. Replacing the audit's individual-mass ratio 0.4596 with the lattice direct ratio $m_u/m_d = 0.474$ (split-Normal, $^{+0.056}_{-0.074}$) shifts the median to ρ = 0.5068, 68% interval [0.480, 0.531], P($\rho \geq 0.500$) = 60% — still fully consistent with ½. Because the direct ratio is the *better-measured* determination (§10), this variant is arguably the more natural headline; the central-value choice for the up/down ratio moves the median by ≈ 0.005, inside the budget, exactly as the linearized analysis anticipated (§10).

The Monte Carlo and the linearized budget agree on the operative statement of the paper: **ρ = 0.50 ± 0.02, with 0.500 essentially central (47th percentile) and the 0.503-vs-0.500 distinction below the present input precision.**

[Established here — Monte Carlo (N = 10^5) confirms ρ = 0.50 ± 0.02 with 0.500 at the 47th percentile (essentially central); ≈ 73% of the variance is m_u/m_d; consistent with the linearized budget.]

Appendix C — QED/Electroweak Differential: a Back-of-Envelope

The audit is QCD/$\overline{MS}$ only (§9). Because χ compares an up-type against a down-type quark, the one effect flavour-universal QCD cancellation does not protect is QED running, whose squared charges differ ($Q^2_{up} = 4/9 \neq 1/9 = Q^2_{down}$). This appendix estimates its size on ρ rather than asserting it negligible. **Crude one-loop, fixed-α, single-reference-scale — order-of-magnitude only.**

The differential. At one loop the electromagnetic mass anomalous dimension is $\gamma_m^{QED} = (3/2)(\alpha/\pi)Q^2$, so the differential on $\chi(g) = \ln(m_{up}/m_{down})$ is

$$d\chi(g) / d\ln\mu\,|_{QED} = -(3/2)(\alpha/\pi)(Q^2_{up} - Q^2_{down}) = -(1/2)(\alpha/\pi) \text{ per unit } \ln\mu,$$

for a same-scale pair. Accumulated from each mass's reference scale up to m_t ($\alpha/\pi \approx 0.0023$):

$\delta\chi\_QED(g) = -(3/2)(\alpha/\pi)[Q^2\_up \ln(m\_t/\mu\_up) - Q^2\_down \ln(m\_t/\mu\_down)]$.

Per-generation shifts (audit reference scales: u,d,s at 2 GeV; c at m_c; b at m_b; t at m_t):

g	pair	δχ_QED(g)
1	u/d (both at 2 GeV)	−0.0051
2	c/s (c at m_c, s at 2 GeV)	−0.0058
3	t/b (t at m_t, b at m_b)	+0.0014

The cancellation, and why it is incomplete. Were all three pairs referenced at a common scale, the δχ_QED(g) would be nearly equal and would cancel *exactly* in ρ (a difference of increments): the idealized "all pairs at 2 GeV" case gives δρ_QED = 0 identically. The audit is not in that case — it references the **top at m_t** (zero QED lever) while the bottom runs up from m_b — and that asymmetry breaks the cancellation:

$\delta\Delta\chi_1 = \delta\chi\_QED(2) - \delta\chi\_QED(1) \approx -0.0007$, $\delta\Delta\chi_2 = \delta\chi\_QED(3) - \delta\chi\_QED(2) \approx +0.0072$,
$\boldsymbol{\delta\rho\_QED \approx (\delta\Delta\chi_2 \cdot \Delta\chi_1 - \Delta\chi_2 \cdot \delta\Delta\chi_1)/\Delta\chi_1^2 \approx +0.0023}$.

Across reasonable mass-definition choices the result ranges from ≈ 0 (idealized common scale) to ≈ +0.0025 (audit scheme, fixed-α range), **positive** in the audit's case — the central +0.0023 uses α/π = 0.0023, and α/π = 0.0025 gives +0.0025. The round headline is **δρ_QED ≈ +0.002**.

Reading. Three conclusions follow. (1) δρ_QED ≈ +0.002 is **subdominant to the ±0.023 input band** — the headline ρ = 0.50 ± 0.02 is unaffected. (2) But it is **comparable to the 0.003 excess above ½ and to the truncation shift, and carries the same (positive) sign** — so part of the central excess is plausibly uncomputed QED, not physics, which is one of the reasons §13 refuses to read the sign of β. There is a corollary worth stating, in the no-information frame and tentatively: since the estimated QED piece is *positive*, subtracting the envelope estimate moves the QCD-plus-structure central value *toward* ½ — from 0.503 to ≈ 0.501 (and the direct-ratio 0.508 to ≈ 0.506) — i.e. possibly onto exact ½. This **strengthens "consistent with β = 0" without reintroducing any sign claim**: the QED estimate is itself crude and ±-uncertain, so this is a tendency, not a correction to be banked, and it does not license reading β in either direction. (3) δρ_QED is therefore **not** negligible at the ρ ≈ 0.001 β-test frontier: a genuine adjacency test needs a proper QED/EW treatment alongside a sharper m_u/m_d. The magnitude and sign depend on the low-scale mass-definition choices (the position of the heavy-quark reference scales), which is itself why a full treatment — not this envelope — is the correct next refinement.

[Estimate — δρ_QED ≈ +0.002 (band 0 to ≈ +0.0025); subdominant to inputs, comparable to the excess and same sign; converts "partly common in the log-ratio structure" from assertion into a bounded, non-vanishing near-cancellation broken by the top-at-m_t reference.]

VERSF Theoretical Physics Programme — Quark Mass Hierarchy Series. Epistemic grade markers (Established / Resolved / Open / Conditional) are applied to truth-apt claims. Gate tags (G-CHI-1 … G-CHI-5) denote standing dependencies. This paper resolves the empirical half of

G-CHI-1 — the QCD/MS̄ scheme-mixing worry is removed and ρ = 0.50 ± 0.02 is consistent with the structurally-predicted ½ and fails to refute it — but does ***not*** *measure exact equality, nor discriminate ½ from nearby alternatives (0.503-vs-0.500 is input-limited, and the central value is itself choice-dependent over [0.503, 0.508]); the value's structural derivation — the Closure-Order Twin and Log-Access Intertwining Lemmas — remains the burden of the structural sequel. Unicode equations throughout; no pole/MS̄ mixing; all intermediate ratios exposed.*