Towards hadronic light-by-light contribution to the muon $g-2$ with C* boundary conditions

Khai Phan Tim Harris Marina Marinkovic Letizia Parato

on behalf of the RC* collaboration

ETH Zürich

Lattice 2026

The muon $g-2$ test

- A sharp precision test of the SM: heavy-physics sensitivity amplified by $(m_\mu/m_e)^2 \approx 4\times 10^4$.
- **Experiment has closed**: E989 final (2025), world average at 124 ppb.
- **Theory has moved**: lattice-driven HVP, $4.2\sigma \to 0.6\sigma$.
- **The test is only as sharp as the theory**: 530 ppb vs. 124 ppb, entirely hadronic.

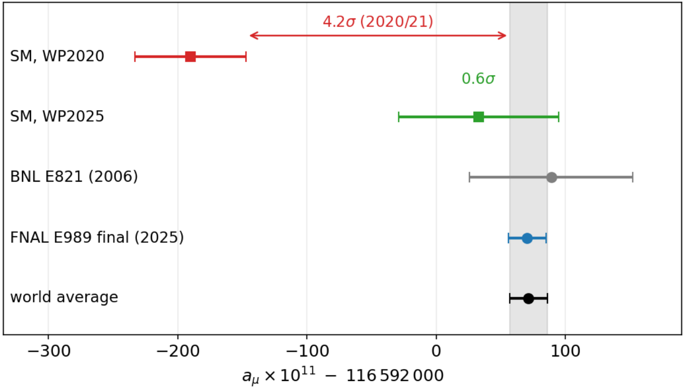

Muon $g-2$ Collab., PRL (2025), arXiv:2506.03069 · Bennett et al. (E821), PRD 73 (2006) 072003, hep-ex/0602035

Aoyama et al. (WP20), Phys. Rept. 887 (2020) 1, arXiv:2006.04822 · Aliberti et al. (WP25), Phys. Rept. 1143 (2025) 1, arXiv:2505.21476

HLbL: the next frontier

- a_μ^{HLbL}: **hadronic light-by-light scattering** dressing the vertex, $\mathcal{O}(\alpha^3)$ (right).
- Largest relative uncertainty of the hadronic budget: 7.4% vs. 0.86% (HVP).
- HVP set to improve further [WP25]: **HLbL becomes the limiting hadronic input**.
- Two mature lattice programmes agree (mild $\approx$1.5σ tension with data-driven); path forward: **independent methods, different systematics**: this talk, a third lattice route.

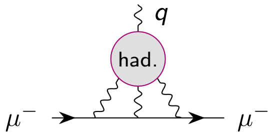

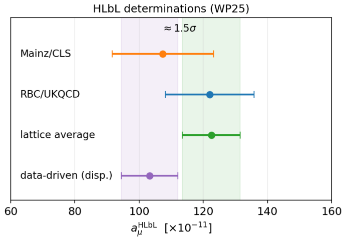

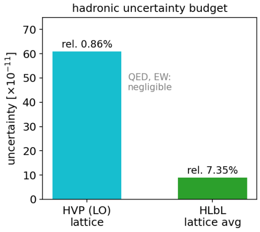

The mainstream route: the four-current correlator

Quantity of interest: the muon vertex dressed by the hadronic four-current correlator,

$$\Pi_{\mu\nu\lambda\rho}(q_1, q_2, q_3) = \int_{x_1 x_2 x_3} e^{-i(q_1 x_1 + q_2 x_2 + q_3 x_3)} \langle J_\mu(x_1) J_\nu(x_2) J_\lambda(x_3) J_\rho(0) \rangle_{\mathrm{QCD}}.$$

Two mature programmes, RBC/UKQCD and Mainz/CLS, compute this **full four-current correlator**; they differ in the QED treatment but share the same computational core.

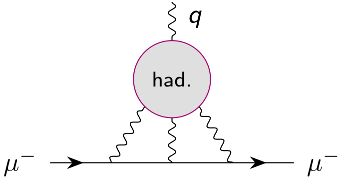

Blum et al. (RBC/UKQCD), PRL 124 (2020) 132002, arXiv:1911.08123 · Chao et al. (Mainz/CLS), Eur. Phys. J. C 81 (2021) 651, arXiv:2104.02632

The subtraction method

$$M = \langle A\,B \rangle_{\mathrm{QCD+QED}} - \langle A \rangle_{\mathrm{QCD+QED}} \langle B \rangle_{\mathrm{QED}}$$

A: loop with two currents (external + internal) B: muon line with one internal current

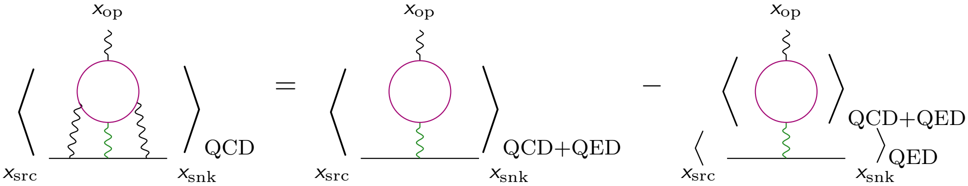

lepton/quark loop muon line photon added by hand (analytic propagator)

- Two two-current correlators on a common ensemble; no four-point function.
- In the difference the lower-order contractions cancel in the path integral; the leading surviving term is the connected HLbL at $\mathcal{O}(\alpha^3)$.

Hayakawa, Blum, Izubuchi, Yamada, PoS LAT2005 (2006) 353, hep-lat/0509016

Blum, Chowdhury, Hayakawa, Izubuchi, PRL 114 (2015) 012001, arXiv:1407.2923

The price and the lever: a controlled cancellation

Signal (mean)

Appears only at $\mathcal{O}(\alpha^3)$: $\langle AB \rangle$ and $\langle A \rangle \langle B \rangle$ are each two orders larger in α; the signal is the survivor of their subtraction.

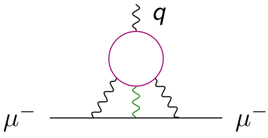

leading connected LbL

Noise (variance)

For $\sigma^2 \sim \alpha^n/N$ the signal-to-noise is $\sqrt{N}\,\alpha^{3-n/2}$; the power n is set by the estimator, not by the observable: noise control is estimator design (quantified near the end).

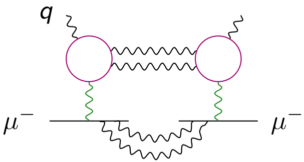

leading variance topology (LLbL, for the estimator used here)

C* boundary conditions, concisely

- Fields are periodic in space **up to charge conjugation**: $\phi(x + \hat{k}L) = \phi^C(x)$.
- The photon is C-odd, so the **zero mode does not exist by construction**: local, gauge-invariant QED in finite volume, with charged states well defined.
- Implemented in **openQxD** by the **RC* collaboration**; dynamical QCD+QED ensembles are part of the RC* programme.

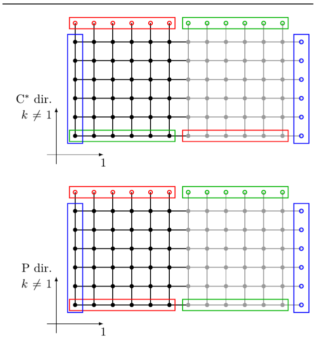

extended (doubled) lattice: physical sites (black), mirror copy (grey)

[openQ*D paper, Fig. 2]

Kronfeld, Wiese, Nucl. Phys. B 357 (1991) 521 · Lucini, Patella, Ramos, Tantalo, JHEP 02 (2016) 076, arXiv:1509.01636

Campos et al. (RC*), Eur. Phys. J. C 80 (2020) 195, arXiv:1908.11673

This project: a three-stage programme

Final objective

Connected + leading disconnected HLbL in fully dynamical QCD+QED (charged sea quarks)

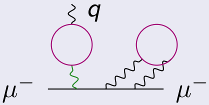

1. QED benchmark

(this talk)

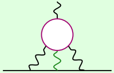

exact perturbative target;
periodic boundaries

$\Rightarrow$

2. QCD + quenched QED

quark loop

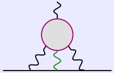

clover Wilson, three spacings,
$M_\pi = 420$ MeV

$\Rightarrow$

3. dynamical QCD+QED

C* ensembles (RC*)

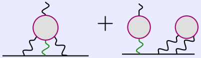

leading disconnected automatic;
ensembles to be generated

Stage 1 tests method, estimator, and systematics where the exact answer is known.

The leptonic test: design, benchmark, and what it is for

- The quark loop is replaced by a lepton loop: exact benchmark (Laporta-Remiddi) $F_2(0)_{\mathrm{LbL}} = 0.371\,(\alpha/\pi)^3$ [Laporta, Remiddi, Phys. Lett. B 301 (1993) 440].
- Boosted coupling $\alpha = 1/4\pi$ ($e = 1$): target $F_2(0) = +6.03 \times 10^{-6}$. Wilson muon, non-compact Feynman-gauge photons; all scales measured on the same ensembles (m_μ, Z_μ, Z_V).
- This talk: coarsest spacing, two geometries:

setup	geometry	N_{cfg}	$\#Q^2$	Q^2_{min}
isotropic	24×8^3	200	4 shells	0.398
anisotropic	$24 \times 16 \times 8^2$	100	10	0.121

The test targets three questions related to the HLbL: **(a)** implementation and cost, **(b)** the continuum limit, **(c)** the statistical nature of the estimator.

From the lattice correlator to a_μ

- The target is the static limit: $a_\mu^{\mathrm{LbL}} = F_2(0)$, the Pauli form factor of the LbL-dressed vertex at zero momentum transfer.
- Projection: the measured three-point tensor, normalized by same-configuration two-point functions, yields $F_2(Q^2)$.

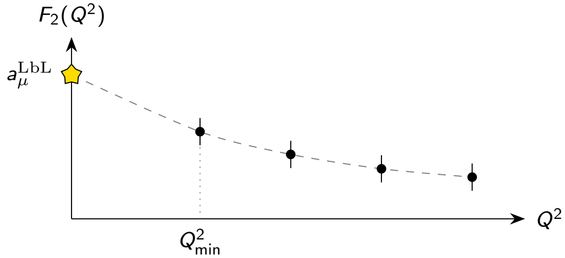

$F_2(Q^2)$ on the isotropic box

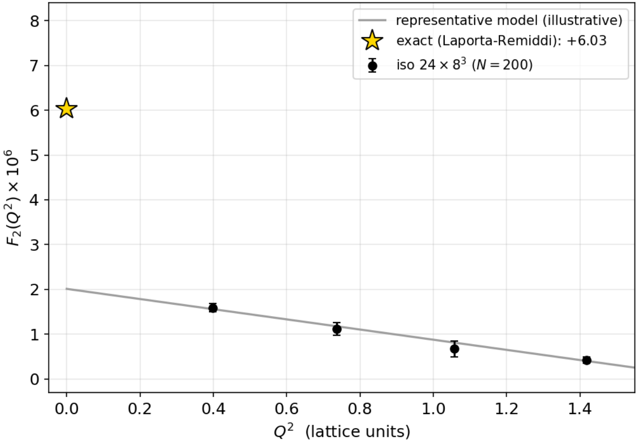

- $N = 200$, four momentum shells, all resolved: signal-to-noise 3.8 to 17.
- Cubic-representation scatter within shells consistent with the statistical errors.
- Positive, monotonically rising toward Laporta-Remiddi as $Q^2 \to 0$ (grey: representative model, illustrative only).
- But: $Q^2_{\text{min}} = 0.398$ leaves a long extrapolation.

Reaching lower Q^2: the anisotropic box

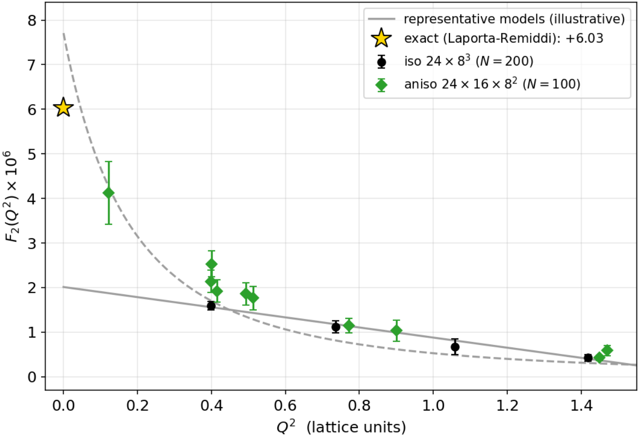

- New lowest point: $F_2(Q^2{=}0.121) = +4.12\,(70) \times 10^{-6}$.
- Grey: representative extrapolation models (illustrative); a controlled $Q^2 \to 0$ extrapolation needs lower Q^2.

Statistical nature: estimator, signal, and variance across the stages

$$\hat{\theta} = \frac{N}{N-1}\left[\frac{1}{N}\sum_{i=1}^{N} A_i^{\sharp}\diamond B_i^{\wedge} - \left(\frac{1}{N}\sum_{i=1}^{N} A_i^{\sharp}\right)\diamond\left(\frac{1}{N}\sum_{j=1}^{N} B_j^{\wedge}\right)\right]$$

	QED	**QCD + qQED**	**dynamical QCD+QED**
estimator	$A_i^{\sharp} = A[\varphi_i]$	$A_i^{\sharp} = A_i^{C}$; $(U_i, \varphi_i) =$ QCD $\times$ quenched U(1)	$A_i^{\sharp} = A_i^{C}$; (U_i, φ_i) sampled jointly
signal	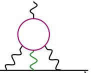	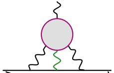	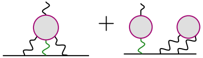
variance	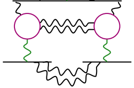	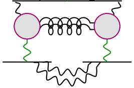	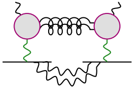

$B_i^{\wedge} = \frac{1}{2}[B_i(e) - B_i(-e)]$: line charge-parity projection, free as the $\pm\mathbf{p}$ average; kills the line's odd-photon components configuration by configuration.

$A_i^{C} = \frac{1}{2}(A[U_i, \varphi_i] + A[U_i^{*}, \varphi_i])$: loop conjugation average, free via the transposition identity on existing solves; restores Furry pointwise on the loop.

The coupling scan: estimator scaling verified in pure QED

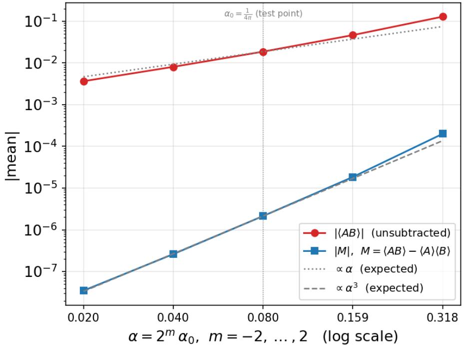

means

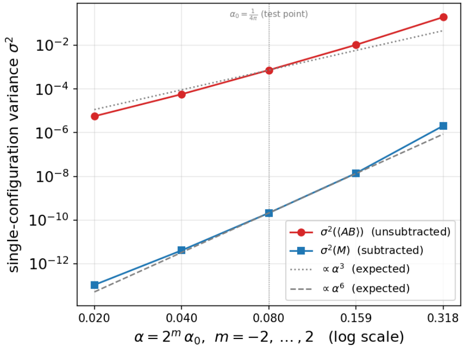

single-configuration variance

- The subtracted mean M follows α^3, as predicted; the unsubtracted $\langle AB \rangle$ is dominated by its $\mathcal{O}(\alpha)$ lower-order piece.
- The subtracted variance follows α^6; the subtraction cancels noise by factors 600 to 8500, so StN $\approx \sqrt{N}$. (Representative momentum; grey lines: expected scaling.)

Next steps: what carries over, what changes for HLbL

QCD + quenched QED

1. Expected variance scaling $\sigma^2 \sim \alpha^4/N$ (quark loops correlate through gluons); to be validated numerically by the same coupling-scan method.
2. Three QCD ensembles available ($M_\pi = 420$ MeV):

a [fm]	lattice	$M_\pi L$
0.049	64×32^3	3.3
0.075	48×24^3	3.8
0.098	32×16^3	3.3

3. Purpose: the continuum-extrapolated fully connected HLbL.

Dynamical QCD+QED

1. Same expected scaling and validation path.
2. Dedicated C* QCD+QED ensembles: to be generated (RC* programme).
3. Automatic inclusion of the leading disconnected HLbL contribution:

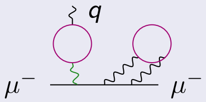

4. **The goal: fully connected + leading disconnected HLbL in fully dynamical QCD+QED with C* boundary conditions.**

Summary

1. The two-current subtraction method is validated in pure QED against the exact Laporta-Remiddi benchmark: $F_2(Q^2)$ resolved (lowest point $F_2(0.121) = +4.12\,(70) \times 10^{-6}$ vs. $+6.03 \times 10^{-6}$).
2. The estimator and its variance are analyzed and validated numerically in QED: signal $\propto \alpha^3$, subtracted variance $\sim \alpha^6$; noise control is estimator design.
3. Next steps: the fully connected HLbL in QCD + quenched QED (three ensembles ready), connected + leading disconnected HLbL in dynamical QCD+QED with C^* boundary conditions.

Computations on CSCS Eiger (Alps).

Backup: QED treatments across programmes

	QED treatment	fermions	FVE strategy
Mainz	QED_∞ kernel	clover Wilson	analytic kernel
RBC/UKQCD	QED_L / QED_∞	domain wall	cross-checked setups
this project	QED_C (C^* BCs)	Wilson	no photon zero mode; FVE under study

Backup: wraparound and time-slice placement

- On a periodic torus each propagator leg of forward length ℓ carries a backward image at $T-\ell$: $A_{\bar{\mu}}/A_{\mu}|_{\mathrm{leg}} = e^{-m_\mu(T-2\ell)}$.
- Leading contamination at a centred insertion: $R_{\mathrm{wrap}} = 2\,e^{-m_\mu(T-t_{\mathrm{sep}})}$.
- Naive placement (source/sink on the temporal boundaries): wrap distance 1, $R_{\mathrm{wrap}} \approx 150\%$, signal buried. This was diagnosed and the layout corrected.
- Corrected symmetric placement on $T=24$: $(t_{\mathrm{src}}, t_{\mathrm{op}}, t_{\mathrm{snk}}) = (7, 12, 17)$, $R_{\mathrm{wrap}} = 2e^{-0.2976\cdot 14} \approx 3\%$.

Backup: sub-sampling benchmark (full table)

M	wall ratio	error ratio (iso)	error ratio (aniso)	cost ratio (iso)
4	0.048	1107	645	5.9×10^4
8	0.057	862	447	4.2×10^4
16	0.073	642	159	3.0×10^4
32	0.101	246	172	6.1×10^3
64	0.168	173	116	5.0×10^3
128	0.298	104	43	3.2×10^3

- Validated against production at $M = V_3$.

Backup: $Q^2 \to 0$ extrapolation

fit	data	$F_2(0) \times 10^6$	class
linear / quadratic	isotropic	$+2.0$ to $+2.3$	polynomial (primary)
quadratic	anisotropic (10 pts)	$+4.0\,(4)$	polynomial
monopole / dipole	isotropic	$+3.1$ / $+7.7$	pole shapes (bracket only)

- QED $F_2(Q^2)$ has a branch cut at $Q^2 = -(2m_\mu)^2$, not a pole: polynomials are the justified fit class.
- Statistical errors on the data are at the few-percent level of the target; the spread above is model dependence plus volume dependence.
- Bottom line: the computation is extrapolation-limited, not statistics-limited.

Backup: extrapolation models (full table)

fit	data	$F_2(0) \times 10^6$	χ^2/dof	class
linear (all 4)	iso	$+2.01\,(12)$	0.49	polynomial (primary)
linear (2 lowest)	iso	$+2.15$	n/a	polynomial (primary)
quadratic	iso	$+2.34\,(37)$	0.11	polynomial (primary)
monopole $M^2{=}(2m)^2$	iso	$+3.11\,(16)$	3.66	pole (bracket only)
dipole $M^2{=}(2m)^2$	iso	$+7.71\,(39)$	3.44	pole (bracket only)
linear	aniso (10 pts)	$+2.60\,(15)$	n/a	polynomial
quadratic	aniso (10 pts)	$+4.01\,(44)$	n/a	polynomial
combined linear	iso + lowest aniso	$+2.07\,(12)$	n/a	mixes volumes
combined quadratic	iso + lowest aniso	$+2.68\,(34)$	n/a	mixes volumes

Backup: normalization and renormalization

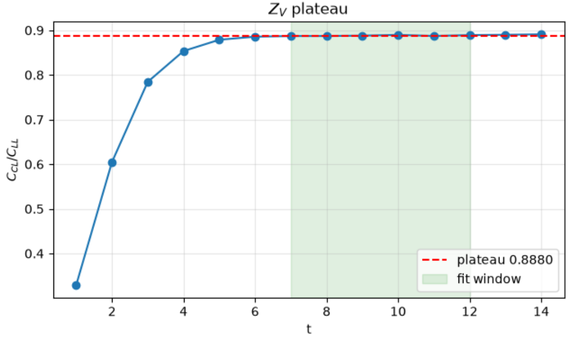

m (muon)	0.2976 (12)	(tuned 0.2976)
Z_μ	0.8041	$(A_0/2)$
Z_V	0.8880 (7)	(plateau $t = 7$-12)
$Z_\mu Z_V^3$	0.5631	
aniso: m, Z_μ	0.2944, 0.7961	

- Fit-window scans flat for m, Z_μ, $E'(q)$.
- Leptonic multiplicity $1/3$ inside F_2^{raw}.

Backup: fit-window stability

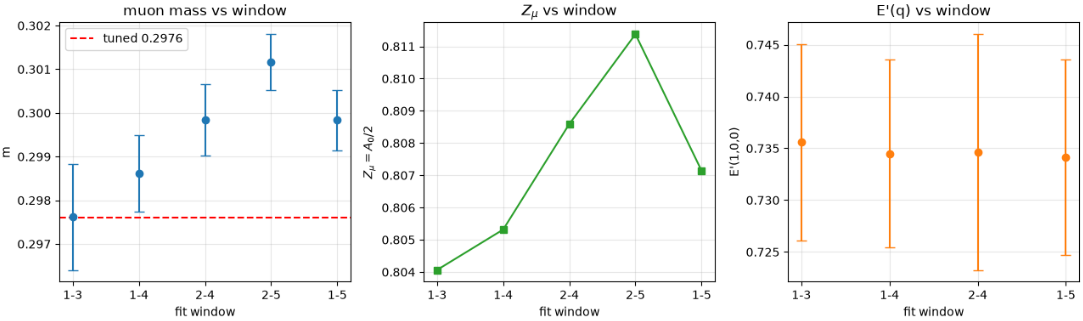

Backup: continuum cost projection

ens	lattice	iso/cfg [node-h]	iso N=100 [node-h]	aniso/cfg [node-h]	aniso N=100 [node-h]
A	24×8^3	0.57 (meas.)	57	1.58 (meas.)	158
B	32×10^3	2.9-3.6	290-360	8-10	800-1000
C	32×12^3	8.7-13	870-1300	24-36	2400-3600
D	40×16^3	61-122	6100-12200	169-337	17000-34000

Backup: implementation and cost: full volume averaging is mandatory

- Wall time is dominated by the $2V_3+1$ propagator solves per configuration: cost $\sim V_3 \cdot V_4$. Measured anchor: 0.57 node-h per configuration on the coarse box.
- The tempting $10\times$ saving: estimate the source loop stochastically on a random subset of M points. Measured verdict: sub-sampling breaks the M_1-M_2 correlation, and the error on the subtracted observable inflates by 10^2 to 10^3 (both geometries, all M; backup table). Cost-to-precision worsens by three orders of magnitude.
- Conclusion, carried to HLbL directly: the loop must be averaged over the full volume; production is budgeted accordingly.

Backup: continuum limit: settled, costed, measurement pending

- Normalization and renormalization under control on the same ensembles: m_μ, Z_μ, $Z_V = 0.8880\,(7)$ measured, overall normalization derived, no external inputs (backup).
- Ensemble chain A-D at fixed $m_\mu L = 2.4$, factor 2 in spacing: generated and tuned (generation costs seconds per configuration in quenched QED).
- Measurement costed from the anchor: B ≈ 300, C $\approx 10^3$, D $\approx 10^4$ node-h at N=100 (isotropic; backup table); single-configuration benchmarks before committing.
- Empirical continuum results: pending; the next production step.

Backup: variance-analysis pipeline and scan method

Variance analysis (the pipeline):

- Expand $\langle \bar{o}^2 \rangle$ over coincidence patterns; three connected structures survive: $\langle ABAB \rangle_c + \langle \widehat{AB} \rangle_c^2 + \langle AA \rangle_c \langle BB \rangle_c$.
- Furry's theorem kills every α^4 term: leading variance α^5/N.
- The $\pm\mathbf{p}$ average ($= e \to -e$ on the muon line, Blum et al. 2015) projects the line onto even photon number: predicted to remove the α^5 term, giving $\sigma^2 \sim \alpha^6/N$ and StN $= \sqrt{N}\,\alpha^0$.

Numerical verification (the trick):

- The coupling lives only in the photon-field normalization: rescaling the stored configurations by $\sqrt{r}$ runs the identical pipeline at $r\alpha$ (validated bit-wise).
- κ retuned per coupling: fixed lepton mass, so fitted slopes are pure powers of α.
- $r = \frac{1}{4} \dots 4$, $N = 50$ each, same Gaussian draws across couplings; ≈ 230 node-h total.

Backup: project status and to-do

	item	status / plan
✓	subtraction method, form factor projection	exact and measured checks
✓	wraparound control (symmetric placement, $T = 3L$)	contamination $\approx 3\%$
✓	estimator study	variance scaling in α verified by scan
✓	leptonic $F_2(Q^2)$, coarsest ensemble, two geometries	$N = 200$ / $N = 100$
☐	leptonic continuum: ensembles B, C, D	costed; benchmarks first
☐	lower Q^2 / volume control	C* momenta, larger L
☐	stage 2: QCD + quenched QED quark loop	3 spacings, $M_\pi = 420$ MeV
☐	stage 3: dynamical QCD+QED on RC* C* ensembles	disconnected automatic

Backup: coupling-scan exponents, full table

quantity	all r	$r \leq 1$	$r \geq 1$	physical power
signal F_2	2.16	1.8	2.4	α^3 at LO
σ^2 subtracted	3.92	3.3	4.7	$\alpha^5 \to \alpha^6$ crossover
σ^2 unsubtracted	1.59	1.25	2.0	$\alpha^3 \to \alpha^4$

- Bookkeeping: physical power $=$ scan slope $+1$ (signal) or $+2$ (variances); the manually inserted photon line does not rescale.
- Local per-octave slopes of σ^2 (subtracted): $3.2, 3.7, 4.2, 5.2$: the physical coupling sits inside the α^5-α^6 crossover.
- Line of constant physics: κ retuned per coupling, $am = 0.297$-0.303 across a factor 16 in α.

Backup: line of constant physics (tuned κ)

r	tuned κ	am	fixed $\kappa = 0.12412$:
1/4	0.117311	0.2973(8)	am drifts 0.13 to 0.91
1/2	0.119356	0.3028(15)	(QED self-energy at
1	0.124120	0.2992(31)	large coupling)
2	0.133848	0.3009(64)	
4	0.155689	0.3033(145)	

- At fixed κ the mass drift injects kinematic exponentials into every raw variance comparison; the tuned chain (secant on the two-point function, seconds per step) removes them, so fitted slopes are pure powers of α. The fixed-κ scan is kept as a systematics check.