A holonomy action for $U(1)$ gauge fields on causal sets: an exact Maxwell mean in 1+1 dimensions and the structure of the higher-dimensional limit

Aaditya Bhatt
Independent Researcher, Ahmedabad 380015, India
ORCID: 0009-0008-4108-5481 aadityasphysics@gmail.com

August 2026

Abstract

We construct a gauge-invariant action for a background $U(1)$ connection on causal sets obtained by Poisson sprinkling of Minkowski spacetime, built from the holonomies of minimal closed causal loops (*chain pairs*) and weighted by Alexandrov-interval cardinality with coefficients of finite support. In 1+1 dimensions we prove that the sprinkling mean of the action on a causal diamond is *exactly* computable at every density ρ, and that with the single coefficient condition $C_3 = -C_2/5$ — which kills the second moment of the pair-count kernel — it converges to the Maxwell action, $\mathbb{E}[S] = \frac{C_2}{5\rho}\int_D F^2\,\mathrm{d}^2x$, for every C^1 field strength, with explicit $O(\rho^{-1/2})$ and boundary corrections. The obstruction is the boost noncompactness of the sprinkling: any nonnegative coefficient family diverges logarithmically, and the canonical Benincasa–Dowker-type alternation $(1,-2,1)$ does *not* cure this action — the chain-pair kernel carries a different combinatorial weight than the Benincasa–Dowker–Glaser layer sums, so its cure is necessarily action-specific. A structural feature accompanies the proof: the same moment condition that cancels the divergence annihilates the entire boost-channel dependence of the smeared holonomy on the field, which is why the theorem survives uniformly over the boost hyperbola. We further prove a variance bound $\mathrm{Var}[S]/\mathbb{E}[S]^2 = O(\ln^2(\rho V_D)/(\rho V_D))$, so the normalized action converges in L^2: the 1+1-dimensional result is an unconditional theorem. In 2+1 dimensions we derive the analogous kill conditions (integer ray $(40,-25,6)$) and the predicted continuum coefficient, prove an exact endpoint-axis degeneracy, and present a numerical pilot consistent with the coefficient at the few-percent level. In 3+1 dimensions we solve the kill system (ray $(14,-16,9,-2)$, log-moment $-\frac{1}{2}$ exactly, scale-free action) and *prove the Maxwell limit for constant field strength on a causal diamond*: a covariogram identity — the Lorentzian analogue of Matheron's covariogram, special to the causal diamond — collapses the mean to an exact double integral, every divergence class carries exactly one killed moment, the surviving class is exactly isotropic (Theorem 8.2; Appendix B, Step 3), and $\mathbb{E}[S] \to \frac{3}{20\pi}\int F_{\mu\nu}F^{\mu\nu}\,\mathrm{d}^4x$, with the bivector constant in closed form ($b_4 = 2/(5\pi)$) and an explicit $\ln\rho\cdot\rho^{-1/2}$ correction series — the conjectured $\rho^{-1/4}$ boundary class does not exist on the diamond. A parity lemma shows further that the kill system and the continuum constant are region-independent, the region entering only through its volume, and the class over which this holds is then *characterized*: a compact causally convex region is exactly a slab $\{f^- \le t \le f^+\}$ between 1-Lipschitz graphs — so cylinders, tubes and slabs are never causally convex — its covariogram obeys an explicit master formula whose first-order term is the Lorentzian analogue of Matheron's perimeter term, an exact linear form in the null coordinates rather than the non-analytic Euclidean $c_1|w|$; the $u \leftrightarrow v$ parity, the absence of a quarter-power class and the constant $\frac{3}{20\pi}\mathrm{Vol}(R)$ all rest on a single regularity condition (M) at null displacements — integer-power expandability of the covariogram at null separations — which also sets the theorem's reach: (M) is proved for the causal diamond, for the polytopal

family and for an exactly solvable smooth family, confirmed numerically on a six-region library, and open for general regions. Null boundary segments — named as the delicate loci in an earlier version — do not break the class. To our knowledge this is the first gauge-field action on a causal set (in the Bombelli–Meyer–Sorkin sense) whose sprinkling mean is proven to converge to the Maxwell action: unconditionally in 1+1 dimensions for C^1 field strengths and in L^2, and in 3+1 dimensions for constant field strength on a causal diamond. The smooth-field analysis in 3+1 dimensions remains open; the *diagonal class* of the 3+1-dimensional variance is shown to *diverge* for this action class, with the divergence law computed exactly (the diagonal class alone gives relative variance $\simeq 42.2\sqrt{\rho V_D}$, and no coefficient ray can cancel it); the $o(\rho^{1/2})$ bound on the signed off-diagonal sharing classes is stated as an open lemma, with numerical evidence that the diagonal class dominates the total (measured total-to-diagonal ratios 0.98–1.22, non-monotone, over the sprinkling-accessible range) — the constant-field result is accordingly to be read as an ensemble-mean theorem — and the one-dimensional escape mechanism is identified. The divergence is nevertheless repairable, but only in a co-scaling window and not at a fixed physical mesoscale: no intrinsically boost-invariant repair by weights or selections built from the configuration geometry exists (the chain-pair configuration has exactly six Lorentz invariants, all boost-blind; bounded functionals of the holonomy itself, such as the compact action $2(1-\cos\Phi)$, are outside that family — for the mean, the leading divergence of its quartic correction is killed by the same moment condition), whereas capping the pair sum at a mesoscale ℓ_m in the region's own rest frame preserves the Maxwell mean *exactly* — the same parity lemma applies with support ℓ_m, and $\ln \ell_m$ multiplies a killed moment — and makes the variance extensive, $\mathrm{Var}[r] \to 3.896\sqrt{\rho}\,\ell_m^6$. Under the cap the repair is an *unconditional* L^2 theorem: a capped interval has region-frame diameter $\leq 2\ell_m$ whatever its boost, which turns the boost logarithm of the add-one estimates into a capped power and bounds every second-moment class, $\mathrm{Var}[S_{\ell_m}] \leq c\,\Lambda_F^4 V_D(\sqrt{\rho}\,\ell_m^6 + \rho\,\ell_m^8)$, while the co-scaled mean error is computed in closed form (the removed tail is an exact finite moment sum whose first survivor is $m_{3/2} = -1.454$); hence a *single* sprinkling's capped killed action converges to the Maxwell action in L^2 along every $\ell_m = \rho^{-a}$ with $\frac{1}{8} < a < \frac{1}{4}$, the wider window $\frac{1}{12} < a < \frac{1}{4}$ retaining the sharp constant conditionally on diagonal dominance. The price stays prominent: at fixed physical ℓ_m the relative variance still diverges as $(\ell_m/\ell_s)^2$, so the cure is a self-averaging window rather than a fixed-scale regulator, and the regulator must vanish in the continuum limit as a lattice spacing does.

1 Introduction

Causal set theory posits that spacetime is fundamentally a locally finite partial order, with the causal relation as the only primitive structure [1, 2]. Its distinctive strength among discrete approaches is Lorentz invariance: a Poisson sprinkling of Minkowski spacetime defines a random causal set that determines no preferred frame, in the sharp sense of the Bombelli–Henson–Sorkin theorem that no equivariant measurable map from sprinklings to directions exists [3]. The price of this invariance is radical nonlocality — every element has infinitely many nearest neighbours along the light cone — and the central technical achievement of the field over the past two decades has been the construction of local continuum physics from this nonlocal discrete structure. For *scalar* fields and for *gravity* the programme is well developed: retarded d'Alembertian operators whose sprinkling means converge to $\Box\phi$ (and, in curved spacetime, to $\Box\phi - \frac{1}{2}R\phi$) were constructed by Sorkin [4] and Benincasa–Dowker [5] in 2 and 4 dimensions and extended to all dimensions [7, 8], yielding the Benincasa–Dowker–Glaser (BDG) action, whose fluctuation and continuum-limit properties have been studied extensively [6, 9, 11]; free scalar quantum field theory on a causal set is available through Johnston's propagator constructions [12, 13].

For *gauge* fields the situation is far less developed. Surya's review states that "at present no well defined representation of non-trivial tensorial fields on causal sets is known" [2], and the recent

Handbook of Quantum Gravity survey of matter on causal sets confirms the same gap [18]. The existing prior art is the corpus of Sverdlov and Bombelli: quasi-local Lagrangians for gravity and matter [17], holonomy-based Lagrangian densities for Yang–Mills fields [14, 15], and an edge-based electromagnetic Lagrangian [16]. These are constructions of candidate actions; to our knowledge none is accompanied by a proof that its sprinkling expectation converges to the Maxwell (or Yang–Mills) action in the continuum limit. (Three adjacent bodies of work bound the claim. A contemporaneous non-arXiv preprint [19] proposes a "diamond holonomy" action on sprinkled causal sets and asserts a 4-dimensional Yang–Mills limit heuristically, resting on an unproven lemma; it contains no convergence theorem, and we distinguish our scoped, theorem-level claims from it. In the *Euclidean* random-lattice setting, gauge theory on Poisson lattices goes back to Christ, Friedberg and Lee [21], likewise without a continuum-limit theorem — which is exactly why the present Lorentzian results do not follow from that literature. And causal *fermion* systems admit a rigorous, book-length continuum limit yielding gauge interactions in a differently-named, non-order-theoretic framework [20]; our claims are scoped to causal sets in the Bombelli–Meyer–Sorkin sense [1].) The present paper supplies convergence theorems in that setting, with the scope stated precisely at the end of the abstract.

We study the simplest gauge-invariant object a causal set offers. Assign to every related pair $p \prec q$ a $U(1)$ phase θ_{pq} (here: the line integral of a background potential along the geodesic). The minimal closed loop built from causal relations alone is a *chain pair*: two elements z, z' in the Alexandrov interval of a pair $x \prec y$, traversed as $x \to z \to y \to z' \to x$. Its holonomy $\Phi = \theta_{xz} + \theta_{zy} - \theta_{xz'} - \theta_{z'y}$ is gauge invariant, and we take as action the cardinality-weighted sum of Φ^2 over all such loops (Section 2). The construction echoes both Wilson's lattice plaquette and the BDG layered-coefficient design, but the analogy is imperfect in an instructive way: the natural "plaquettes" of a sprinkling form a Lorentz-invariant ensemble of unbounded coordinate size, and the sum over them diverges logarithmically — the same boost noncompactness that forces the BDG alternating-coefficient structure [4, 5]. Our main results are:

1. **An exact mean at every density** (Theorem 3.2). In 1+1 dimensions, for constant field strength on a causal diamond, $\mathbb{E}[S]$ is given in closed form at every ρ — including all boundary corrections. The raw (positive-coefficient) scheme diverges as $\ln \rho$ with computable constant; the moment-kill condition $C_3 = -C_2/5$ yields $\mathbb{E}[S] \to \frac{C_2}{5} E^2 V_D/\rho$. Neither nonnegative coefficients nor the canonical BDG-type alternation can achieve this (Remark 3.4); the condition is action-specific.

2. **The Maxwell limit for smooth fields** (Theorem 4.1). For every C^1 field strength the mean action localizes on the Maxwell density: $\rho\,\mathbb{E}[S] \to \frac{C_2}{5} \int_D F^2$. The proof turns on Lemma 4.4: after an exact reduction of the pair measure to volume and boost variables, *every* contribution of the field profile along the boost (null-stretch) channel is annihilated by the same moment condition that kills the divergence. The divergence cure and the uniform-in-boost remainder control are one and the same condition — the structural reason a Lorentz-invariant discrete gauge action can exist at all.

3. **Variance control, hence an unconditional theorem** (Theorem 5.3). Via the Poincaré inequality for Poisson functionals we prove $\mathrm{Var}[S]/\mathbb{E}[S]^2 = O(\ln^2(\rho V_D)/(\rho V_D))$, so the normalized action converges to the Maxwell action in L^2, not merely in mean. The proof requires a second-moment bound on the add-one-point cost (Lemma 5.2, proved in Appendix A), and explains quantitatively the anomalously slow $\rho^{-1/3}$ self-averaging observed in simulations as the finite-density face of a $\ln \rho \cdot \rho^{-1/2}$ law.

4. **The higher-dimensional structure** (Sections 7–8). In $d > 2$ the loop bivector is no longer boost invariant and the divergences are power law in the boost rapidity, organized — thanks to an exact endpoint-axis degeneracy (Proposition 7.1) — into finitely many moment classes. In 2+1 dimensions the resulting kill system has the integer solution $(C_2, C_3, C_4) = (40, -25, 6)$ and the exact reduced mean integral approaches the predicted continuum coefficient $-\frac{16}{3} b_3 m_{4/3} = -2.7647$ — a pilot consistent at the few-percent level, not yet a confirmation. In 3+1 dimensions we solve the kill system $(C_2, \ldots, C_5) = (14, -16, 9, -2)$ with surviving log-moment exactly $-\frac{1}{2}$, observe that the resulting action is scale-free — Maxwell marginality emerges from the counting — and then prove the 3+1-dimensional Maxwell limit for constant field strength on a causal diamond (Theorem 8.2): the loop bivector is exactly $\Sigma^{0i} = \frac{\tau}{2}(z' - z)^i$, giving $b_4 = 2/(5\pi)$ in closed form; a covariogram identity ($D \cap (D - w)$ is itself an Alexandrov interval) collapses the region integral exactly; and the resulting closed-form mean converges to $\frac{3}{20\pi} \int F_{\mu\nu} F^{\mu\nu}$ with an explicit correction series (Appendix B). The surviving divergence class is exactly isotropic — a consequence of the exact region integral, whose covariogram is independent of the spatial axis direction (Remark 8.1 recapitulates the mechanism heuristically). A parity lemma (Lemma 8.4) proves the kill system and the continuum constant region-independent — the region enters only through its volume — and the class over which that holds is then *characterized* (Lemma 8.5, Proposition 8.6): a compact causally convex region is exactly a 1-Lipschitz slab $\{f^- \le t \le f^+\}$, so cylinders and slabs are not hard cases but non-cases; its covariogram's first-order Matheron perimeter term is an exact linear form in (u, v), causal convexity being what analytises the Euclidean $c_1|w|$; the parity, the absence of quarter-power classes and the constant all rest on a single null-displacement regularity condition (M), which also sets the reach of the theorem: proved for diamonds, for the polytopal family and for the smooth family, open in general. Null boundary segments do not break the class. What remains — that condition in complete generality — and the smooth-field 3+1-dimensional extension are posed precisely (Open Problem 8.8).

5. **The diagonal class of the four-dimensional variance diverges** (Section 9). The variance question that Theorem 5.3 settles affirmatively in 1+1 dimensions is answered *negatively* in 3+1 at the level of the diagonal class: the diagonal class of the second Mecke moment is exactly computable, grows as $\rho^{1/2}$ (relative variance $\simeq 42.2\sqrt{\rho V_D}$), and — because its kernel carries the *squares* C_n^2 with strictly positive weights — cannot be cancelled by any real finite-support coefficient family (Proposition 9.2). The $o(\rho^{1/2})$ bound on the signed off-diagonal sharing classes is *not* proved (Section 9.2; Open Problem 9.10(a)); the numerical evidence is that the diagonal class dominates the total. The 1+1-dimensional escape is identified: there the loop observable is boost invariant (Lemma 3.1), so the boost tail that powers the divergence is absent. Theorem 8.2 is accordingly to be read as an *ensemble-mean* theorem — an interpretive choice justified by Theorem 9.1 together with the numerical evidence of Section 9.2, not by a proved two-sided divergence — while the 1+1-dimensional L^2 theorem is unaffected.

6. **What repairs the variance, and what it costs** (Section 9.4). The divergence turns out to be a property of the observable, not of the substrate — but no *intrinsic* repair by configuration-geometry weights exists: the chain-pair configuration carries exactly six Lorentz invariants, all of them rest-frame quantities, so no invariant weight or selection built from them can see the boost (Proposition 9.4; bounded per-loop functionals of the holonomy itself, such as the compact action, are outside that family — their mean is controlled at quartic order by Lemma 9.11, their variance open); and invariant-norm plaquettes do insert the hoped-for sinh

factors, at identical leading coefficient (Proposition 9.5). What works is a region-covariant mesoscale cap: restricting the pair sum to null extents $\max(u, v) \le \ell_m$ in the region's own rest frame preserves the Maxwell mean *exactly* — Lemma 8.4 applies verbatim with support ℓ_m, and the kill that makes the mean finite makes it cap-independent — while restoring extensivity of the variance (Proposition 9.7). The price is stated as loudly as the cure: at fixed physical ℓ_m the relative variance still diverges, so what is recovered is a self-averaging *window* $\ell_s \ll \ell_m \ll 0.364\rho^{-1/12}T^{2/3}$, with the regulator vanishing in the continuum limit as it does on a lattice. Under the cap the repair is complete as an L^2 statement: all second-moment classes are bounded unconditionally (Proposition 9.8), the co-scaled mean error is computed in closed form (22), and the region-dependence question of Open Problem 8.8(b,c) disappears — the cap leaves the region in the answer only through its volume.

Every numerical claim in the paper is reproduced, and the frozen ones asserted, by an ancillary script (Section 6). We work throughout with a fixed background (classical) connection; quantization — for which the indefiniteness of the moment-killed action is the salient issue, precisely as for the BDG gravitational action [27, 28] — and the non-abelian extension are discussed, with sharply posed first computations, in Section 10.

2 The action

2.1 Sprinklings, phases, loops

Let D be a causally convex compact region of d-dimensional Minkowski spacetime $\mathbb{M}_d$ of volume V_D, and let η be a Poisson point process on D of intensity ρ (a *sprinkling*); write ℓ for the discreteness length, $\rho = \ell^{-d}$. For $x \prec y$ (causal order) let $I(x, y)$ denote the open Alexandrov interval $I^+(x) \cap I^-(y)$, V_{xy} its volume, and $n(x, y) = |\eta \cap I(x, y)|$ its cardinality (endpoints excluded). Causal convexity of D guarantees $I(x, y) \subset D$ for all $x, y \in D$.

Fix a smooth $U(1)$ potential A_μ on D and assign to each related pair the phase

$$\theta_{pq} = \int_p^q A_\mu \, \mathrm{d}x^\mu \qquad \text{(straight segment from } p \text{ to } q\text{)}. \tag{1}$$

Under a gauge transformation $A \to A + \mathrm{d}\lambda$ one has $\theta_{pq} \to \theta_{pq} + \lambda(q) - \lambda(p)$, so any *closed* combination of phases is gauge invariant. The minimal closed structure available from causal relations is the *chain pair*: an unordered pair $\{z, z'\} \subset \eta \cap I(x, y)$, traversed as the loop $x \to z \to y \to z' \to x$, with holonomy phase

$$\Phi_{xy}(z, z') = \theta_{xz} + \theta_{zy} - \theta_{xz'} - \theta_{z'y}. \tag{2}$$

By Stokes' theorem $\Phi = \int_Q F$, the flux of $F = \mathrm{d}A$ through the (signed) quadrilateral Q with vertices x, z, y, z'; in particular $|\Phi| \le 2\,\Lambda_F V_{xy}$ with $\Lambda_F = \sup_D |F|$, since Q lies in the convex region $I(x, y)$.

2.2 The action and its choices

The action is

$$S \;=\; \sum_{x \prec y \in \eta} C_{n(x,y)} \sum_{\{z,z'\} \subset \eta \cap I(x,y)} \Phi_{xy}(z, z')^2, \tag{3}$$

with real coefficients C_n of *finite support* in $n \ge 2$. The design choices, stated so that they can be disagreed with: (i) *all* causal pairs (x, y) enter, weighted by interval cardinality — not only

links; the finite-support cardinality weight is what makes the all-pairs sum convergent, exactly as in the BDG layered design [5, 7]; (ii) intermediates z, z' range over all interval elements; (iii) the local loop functional is Φ^2, the small-holonomy limit of the compact choice $2(1-\cos\Phi)$; Φ^2 is the one we control analytically, and the relation between the two is *not* innocuous in 3+1 dimensions — the variance divergence of Section 9 is carried by configurations with $\Phi^2 \to \infty$, exactly where the bounded compact functional and the quadratic one differ maximally. At the level of the *mean*, however, the leading divergence class of the quartic correction $\mathbb{E}[\Phi^4]/12$ is annihilated by the same moment-kill condition (Lemma 9.11), so the constant-field mean theorem survives compactification at quartic order; sextic and higher orders are not accounted, and every variance statement in this paper is proved for Φ^2 only; (iv) phases are assigned by the background-field prescription (1); dynamical link phases are beyond our scope (Section 10). All results below concern the random variable (3) under the sprinkling measure.

A dimensional remark orients the whole paper. The kernel weight concentrates the sum at interval volumes of order the discreteness scale, $V \sim \rho^{-1}$, where a loop of linear size $V^{1/d}$ carries $\Phi^2 \sim F^2 V^{4/d}$; two Mecke integrations against the exponentially weighted kernel then give $\mathbb{E}[S] \sim \rho^2 \cdot \rho^{-1-4/d} \cdot V_D F^2 = \ell^{4-d} \int F^2$: an irrelevant coupling in $d = 2, 3$ and a *marginal* one in $d = 4$ — the causal-counting action reproduces the scale-free character of 4-dimensional Maxwell theory automatically (Section 8).

3 The exact mean in 1+1 dimensions

Work in lightcone coordinates $u = t - x$, $v = t + x$, in which D is the box $[0, L]^2$ (a causal diamond of volume $V_D = L^2/2$; the spacetime measure is $\mathrm{d}t\,\mathrm{d}x = \frac{1}{2}\mathrm{d}u\,\mathrm{d}v$), and let the field strength be constant, $F_{tx} = E$. Define the *pair-count kernel*

$$k(s) \;=\; e^{-s} \sum_n C_n \frac{s^n}{n!} \frac{n(n-1)}{2}, \qquad m_\alpha = \int_0^\infty s^\alpha k(s)\,\mathrm{d}s, \quad m_\alpha^{\log} = \int_0^\infty s^\alpha k(s) \ln s\,\mathrm{d}s. \tag{4}$$

Explicitly $m_\alpha = \sum_n C_n \frac{n(n-1)}{2} \frac{\Gamma(n+1+\alpha)}{n!}$; for support $\{2, 3\}$, $m_2 = 12C_2 + 60C_3$.

Lemma 3.1 (interior holonomy, exact)**.** *Let $x \prec y$ span an interval of volume V (u-, v-side lengths u_0, v_0, $V = u_0 v_0/2$) and let z, z' be i.i.d. uniform in $I(x, y)$. For constant $F = E$,*

$$\Phi = E\,A, \qquad A = \tfrac{1}{4}\big[u_0(v_2 - v_1) - v_0(u_2 - u_1)\big], \qquad \mathbb{E}[A^2] = \frac{V^2}{12},$$

where (u_i, v_i) are the lightcone coordinates of z, z' relative to x. Moreover A is invariant under boosts ($u \to e^{\chi} u$, $v \to e^{-\chi} v$).

Proof. For two points p, q, the (t, x)-plane cross term is $t_p x_q - t_q x_p = \frac{1}{2}(u_p v_q - u_q v_p)$, so the shoelace area of the loop $x \to z \to y \to z' \to x$ with x at the origin and $y = (u_0, v_0)$ reduces to the displayed A. The cross moment $\mathbb{E}[(v_2 - v_1)(u_2 - u_1)]$ vanishes by independence of the u- and v-coordinates of a uniform point in the box, and $\mathbb{E}[(v_2 - v_1)^2] = v_0^2/6$, giving $\mathbb{E}[A^2] = \frac{1}{16}\big[u_0^2 \frac{v_0^2}{6} + v_0^2 \frac{u_0^2}{6}\big] = u_0^2 v_0^2/48 = V^2/12$. Boost invariance is manifest. □

Theorem 3.2 (exact mean, constant F)**.** *For any finite-support coefficients, on the diamond D_L at density ρ, with $w = 2V$ the lightcone-box area of an interval,*

$$\mathbb{E}[S] \;=\; \frac{\rho^2 E^2}{192} \int_0^{L^2} w^2\, k\big(\tfrac{\rho w}{2}\big) \Big[(L^2 + w) \ln \tfrac{L^2}{w} - 2(L^2 - w)\Big] \mathrm{d}w, \tag{5}$$

exactly. Consequently, with $T_0 = \rho L^2/2 = \rho V_D$ *and* $r := \rho\,\mathbb{E}[S]/(E^2 V_D)$,

$$r(\rho) = \frac{1}{12}\Big\{m_2\big(\ln T_0 - 2\big) - m_2^{\log} + \frac{1}{T_0}\Big[m_3\big(\ln T_0 + 2\big) - m_3^{\log}\Big]\Big\} + O\big(e^{-T_0} T_0^{n_{\max}+3}\big). \tag{6}$$

Proof. By the multivariate Mecke equation [26], summing over the pair (x, y) and conditioning on $n(x, y) = n$ (given which the interval points are i.i.d. uniform),

$$\mathbb{E}[S] = \rho^2 \iint_{x \prec y \in D} \mathrm{d}x\,\mathrm{d}y\; \frac{E^2 V_{xy}^2}{12}\; k(\rho V_{xy}),$$

using Lemma 3.1 and $\sum_n C_n\, e^{-\rho V}\frac{(\rho V)^n}{n!}\frac{n(n-1)}{2} = k(\rho V)$. Parametrizing $y - x = (a, b)$ in lightcone coordinates, $\iint \mathrm{d}x\,\mathrm{d}y\, g = \frac{1}{4}\int_0^L \mathrm{d}a \int_0^L \mathrm{d}b\,(L-a)(L-b)\, g(ab/2)$, and the substitution $w = ab$ with the exact inner integral

$$\int_{w/L}^{L} \frac{\mathrm{d}a}{a}(L-a)\Big(L - \frac{w}{a}\Big) = (L^2 + w)\ln\frac{L^2}{w} - 2(L^2 - w)$$

gives (5). Substituting $s = \rho w/2$ and extending the truncated moments $\int_0^{T_0} \to \int_0^\infty$ (error $O(e^{-T_0} T_0^{n_{\max}+3})$ from the exponential kernel) yields (6). □

The logarithm in (5) is the *boost noncompactness* of the sprinkling: intervals of fixed volume V form a hyperbola of unbounded coordinate extent, cut off only by the region — the same structure that governs the BDG operators [4, 5, 6]. Since the summand Φ^2 is positive and (by Lemma 3.1) boost invariant, no coefficient choice can cancel the divergence pointwise on the hyperbola; only the *moment* m_2 can be killed.

Corollary 3.3 (divergence and cure)**.** *(i) For* $C_n \geq 0$ *(not all zero),* $m_2 > 0$ *and* $r(\rho) = \frac{m_2}{12}\ln\rho + O(1)$ *diverges logarithmically. (ii) For support* $\{2, 3\}$ *the unique cure is the* moment-kill ray

$$m_2 = 0 \iff C_3 = -\tfrac{C_2}{5}, \qquad \textit{giving} \qquad m_2^{\log} = -\tfrac{12}{5}C_2 \tag{7}$$

(from $\psi(6) - \psi(5) = \frac{1}{5}$*), and then, exactly up to the exponential tail,*

$$r(\rho) = \frac{C_2}{5} - \frac{C_2}{T_0}\Big(\ln T_0 + 1 - \psi(6)\Big), \qquad \mathbb{E}[S] \xrightarrow[\rho\to\infty]{} \frac{C_2}{5}\,\frac{E^2 V_D}{\rho} = \frac{C_2\,\ell^2}{5}\int_D F^2. \tag{8}$$

(iii) For the raw scheme (C_2 *only),* $r_{\mathrm{raw}}(\rho) = C_2\big[\ln T_0 - 2 - \psi(5) + \frac{5}{T_0}(\ln T_0 + 2 - \psi(6))\big]$.

Proof. All statements follow from (6) by evaluating the moments: $m_2^{\log} = \sum_n C_n \frac{n(n-1)}{2}\frac{\Gamma(n+3)}{n!}\psi(n+3)$, and for support $\{2, 3\}$ on the kill ray, $m_3 = -12C_2$, $m_3^{\log} = m_3(\psi(6) + 1)$, which collapses the boundary bracket in (6) to the displayed form. □

The finite-density formula (8) is not merely an error bound: it reproduces the simulated means at all accessible densities within Monte Carlo errors, including the visible deficit at $\rho V_D = 125$ (Section 6, Figure 1).

Remark 3.4 (the canonical alternation does not cure this action). The BDG scalar constructions are built on alternating layer coefficients of the type $(1, -2, 1)$ [5, 7]. Transplanted to the pair-count kernel (4) (support $\{2, 3, 4\}$), they give $m_2 = 12 - 120 + 180 = 72 \neq 0$ (placed at $n = 0, 1, 2$ instead, $m_2 = 12 \neq 0$): the action still diverges. The structural reason should be stated plainly,

lest a deeper obstruction be inferred: the pair-count kernel (4) carries the combinatorial weight $\binom{n}{2}$ over the elements of a single interval, whereas the BDG smearing weights are $\binom{n}{k}$-structured over *layers* of the interval cardinality of an element pair [5, 7] — the extra factor $n(n-1)/2$ shifts the effective Gamma argument by two, so different coefficients are forced by construction, before any dynamics enters. The correct reading is not that the canonical cure mysteriously fails, but that the loop observable carries a different combinatorial weight, for which the BDG coefficients are simply not the relevant ones; the condition (7) is the one relevant here, and had to be derived for this action. The moment-kill *technique* itself, we should say plainly, is not new: solving a linear system of moment-cancellation conditions on alternating-sign layer coefficients, so that the faster-growing terms vanish while the desired leading term is normalized, is the established cancellation mechanism of the BDG/Sorkin family, stated in general form by Aslanbeigi, Saravani and Sorkin [10]; what is new here is the observable it is applied to (chain-pair holonomies) and the specific systems it must solve, (7), (10) and (12). Alternation itself, however, is *forced*: by Corollary 3.3(i), positivity of the coefficients is incompatible with a finite continuum limit.

4 Smooth fields: the boost-channel annihilation

Let now $F = F_{tx}(t,x)$ be an arbitrary C^1 field strength on $\bar{D}$. The result of this section is:

Theorem 4.1 (Maxwell mean, smooth F)**.** *For the moment-killed coefficients* (7)*,*

$$\mathbb{E}[S] \;=\; \frac{C_2}{5\rho}\int_D F^2\,\mathrm{d}t\,\mathrm{d}x\,\Big[1+O\big(\rho^{-1/2}\big)+O\big(\ln(\rho V_D)/(\rho V_D)\big)\Big], \tag{9}$$

with constants depending on $\|F\|_\infty$, $\|\partial F\|_\infty$, L.

The danger is uniformity over the boost hyperbola: an interval of tiny volume can be enormously stretched along a null direction, so its loops sample F over $O(L)$ coordinate distances, and naive Taylor bounds fail. The proof turns this danger into the theorem. Parametrize an interval by its center $\bar{x}$, volume V, and side lengths u_0, v_0; by the reflection symmetry $u \leftrightarrow v$ assume $u_0 \geq v_0$ (long direction u) and define the *boost variables*

$$p = \tfrac{1}{2}u_0^2 \in [V,\, p_{\max}(\bar{x})], \qquad q = \tfrac{1}{2}v_0^2 \in (0, V], \qquad pq = V^2.$$

The exact pair measure factorizes as $\mathrm{d}x\,\mathrm{d}y = \mathrm{d}\bar{x}\,\mathrm{d}V\,\frac{\mathrm{d}p}{p}$ on this chart: the naive Jacobian of the unordered chart gives $\frac{1}{2}\,\mathrm{d}V\,\mathrm{d}p/p$, and the $u \leftrightarrow v$ reduction to the ordered half $u_0 \geq v_0$ restores $\mathrm{d}V\,\mathrm{d}p/p$ (Jacobian computation as in Theorem 3.2); $p_{\max}$ is V-independent.

Lemma 4.2 (central symmetry)**.** *Fix $x \prec y$, let $\bar{x} = \frac{1}{2}(x+y)$ be the interval centre, let $\sigma(z) = x + y - z$ be the point reflection through $\bar{x}$, an involution of $I(x,y)$ preserving Lebesgue measure, and let $\Delta(\xi) = \xi - \bar{x}$ denote the displacement of a point ξ from the centre. Write $\Phi = F(\bar{x})A + \mathcal{T} + O(\|\Delta\|^2)$, where A is the signed area of Lemma 3.1 and $\mathcal{T} = \int_Q (\Delta\cdot\nabla F)(\bar{x})$ is the first-order Taylor functional (the flux of the linearized field through Q). Applying σ to both intermediates,*

$$A \circ \sigma = -A, \qquad \mathcal{T} \circ \sigma = +\mathcal{T},$$

so $\frac{1}{2}\big[\Phi^2 + (\Phi\circ\sigma)^2\big] = F(\bar{x})^2A^2 + \mathcal{T}^2 + O(\|\Delta\|^3)$: the first-order cross term $2F(\bar{x})A\mathcal{T}$ vanishes under the (z,z') average, and the first-order Taylor term of $\mathbb{E}[\Phi^2]$ vanishes identically at every $(V,p,\bar{x})$.

Proof. Represent flux through the closed loop by its winding function: for ξ off the loop let $n_Q(\xi) \in \{-1, 0, 1\}$ be the winding number of $x \to z \to y \to z' \to x$ around ξ (a closed quadrilateral winds at most once), so that $\int_Q \omega = \int n_Q\, \omega$ for every 2-form ω; in particular $A = \int n_Q$ and $\mathcal{T} = \int n_Q(\xi)\, (\Delta(\xi) \cdot \nabla F)(\bar{x})\, \mathrm{d}\xi$. The reflected loop $x \to \sigma(z) \to y \to \sigma(z') \to x$ is the σ-image of the original loop traversed in the reverse order: σ maps the vertex cycle (x, z, y, z') to $(y, \sigma z, x, \sigma z')$, which is cyclically the reversed cycle $(x, \sigma z', y, \sigma z)$. A point reflection of the plane is a rotation by π — orientation preserving — while reversal of traversal negates the winding, so

$$n_{Q\circ\sigma}(\xi) \;=\; -\, n_Q\big(\sigma(\xi)\big).$$

The two signs now follow by the substitution $\xi = \sigma(\eta)$ (unit Jacobian; $\Delta \circ \sigma = -\Delta$):

$$A \circ \sigma = \int n_{Q\circ\sigma} = -\int n_Q(\sigma(\xi))\, \mathrm{d}\xi = -A, \qquad \mathcal{T} \circ \sigma = -\int n_Q(\eta)\big(-\Delta(\eta) \cdot \nabla F\big)(\bar{x})\, \mathrm{d}\eta = +\mathcal{T}.$$

For A the sign is also visible directly from the shoelace formula: with x at the origin, $\sigma(z) = y - z$ sends $(u_i, v_i) \mapsto (u_0 - u_i,\, v_0 - v_i)$, and Lemma 3.1 gives $A \mapsto \frac{1}{4}\big[u_0\big((v_0 - v_2) - (v_0 - v_1)\big) - v_0\big((u_0 - u_2) - (u_0 - u_1)\big)\big] = -A$. Since σ preserves the uniform measure on $I(x, y)$ and z, z' are i.i.d. uniform, the (z, z') average of Φ^2 equals that of $\frac{1}{2}[\Phi^2 + (\Phi \circ \sigma)^2]$, in which the cross term is absent; the $O(\|\Delta\|^3)$ remainder is controlled by the C^1 modulus as in Lemma 4.3. (Note that composition with $z \leftrightarrow z'$ is not needed — the reflection of the intermediates alone suffices.) □

Lemma 4.3 (null-strip decomposition)**.** *Write $\Phi = \int h_Q(u)\, F(u, \bar{v})\, \mathrm{d}u + r$, where the signed width profile of Q at u is $h_Q(u) = \int n_Q(u, v)\, \mathrm{d}v$, with n_Q the winding function of Lemma 4.2 and normalization $\int h_Q\, \mathrm{d}u = A$ (the signed area of Lemma 3.1), and the line $v = \bar{v}$ passes through the interval center. Then $|r| \le 2V \|\partial_v F\|_\infty \sqrt{2q}$, and*

$$\mathbb{E}\Big[\Big(\int h_Q F\, \mathrm{d}u\Big)^2\Big] = V^2\, \delta_\parallel(p; \bar{x}),$$

an exact function of $(p, \bar{x})$ alone (the v-dependence of F having been discarded into r), with $\delta_\parallel(0^+; \bar{x}) = F(\bar{x})^2/12$ and

$$\big|\delta_\parallel(p; \bar{x}) - \delta_\parallel(0^+; \bar{x})\big| \le \sqrt{2}\, \|\partial F\|_\infty \big(4\|F\|_\infty + L\|\partial F\|_\infty\big)\sqrt{p}.$$

Proof. Step 1: the shape law. The structural ingredient is that the *shape law* of the normalized width profile is independent of the boost parameter p at fixed V; we prove it by exhibiting the boost map explicitly. At fixed V, the linear map $\beta_\lambda : (u, v) \mapsto (\lambda u, \lambda^{-1} v)$ $(\lambda > 0)$ carries the interval box $[0, u_0] \times [0, v_0]$ onto the interval box $[0, \lambda u_0] \times [0, \lambda^{-1} v_0]$ of the boosted pair with the same volume; having unit Jacobian, it pushes the uniform law on the first box to the uniform law on the second, so the configuration (z, z') of the boosted interval is β_λ applied to a configuration of the unboosted one, in law. The width profile is equivariant under this map: the quadrilateral of the boosted configuration is $\beta_\lambda Q$, and since β_λ stretches the u-axis by λ and compresses widths (which are v-lengths) by λ^{-1},

$$h_{\beta_\lambda Q}(\lambda u) = \lambda^{-1}\, h_Q(u).$$

Consequently the dimensionless profile $g(s) := (u_0/V)\, h_Q(u_0 s)$, $s \in [0, 1]$, has a law depending only on the shape of the box — i.e. on nothing, at fixed V, since any two boxes of equal volume are related by some β_λ. Writing the strip integral in terms of g,

$$\int h_Q(u)\, F(u, \bar{v})\, \mathrm{d}u = V \int_0^1 g(s)\, F(u_0 s, \bar{v})\, \mathrm{d}s,$$

whose second moment is V^2 times a functional of the law of g and of F along a strip of length $u_0 = \sqrt{2p}$ through $\bar{x}$ — an exact function of $(p, \bar{x})$ alone, as claimed (all coordinates relative to the interval, the strip centred at $\bar{x}$); the exactness is that of the truncated functional $\int h_Q F(u, \bar{v})\,\mathrm{d}u$, whose discarded v-dependence is precisely r, bounded in Step 2.

Step 2: the remainder bound. By Stokes' theorem and the winding representation, $\Phi = \iint n_Q(u,v)\,F(u,v)\,\mathrm{d}u\,\mathrm{d}v$, and by the definition of h_Q,

$$r = \iint n_Q(u,v)\,\big[F(u,v) - F(u,\bar{v})\big]\,\mathrm{d}u\,\mathrm{d}v.$$

On the support of n_Q — contained in the interval box — one has $|v - \bar{v}| \le v_0/2 = \sqrt{q/2}$, hence $|F(u,v) - F(u,\bar{v})| \le \|\partial_v F\|_\infty \sqrt{q/2}$; since $|n_Q| \le 1$ and the box has area $u_0 v_0 = 2V$,

$$|r| \;\le\; 2V\,\|\partial_v F\|_\infty \sqrt{q/2} \;\le\; 2V\,\|\partial_v F\|_\infty \sqrt{2q}\,.$$

Step 3: the $\sqrt{p}$ modulus. Write $F(u_0 s, \bar{v}) = F(\bar{x}) + e(s)$ with, along the strip, $|e(s)| \le \|\partial F\|_\infty\,|u_0 s - \frac{u_0}{2}| \le \|\partial F\|_\infty \sqrt{p/2}$. Two deterministic bounds: the section of Q over u lies in $[0, v_0]$, so $|h_Q(u)| \le v_0$ and $\int_0^1 |g(s)|\,\mathrm{d}s = \frac{1}{V}\int |h_Q(u)|\,\mathrm{d}u \le \frac{u_0 v_0}{V} = 2$; consequently also $|\int_0^1 g| = |A|/V \le 2$ and $\big|\int_0^1 g\,e\big| \le 2\|\partial F\|_\infty\sqrt{p/2} = \sqrt{2}\,\|\partial F\|_\infty\sqrt{p}$. As $p \to 0$ the strip contracts to the center and $\delta_\parallel(0^+;\bar{x}) = F(\bar{x})^2\,\mathbb{E}[(\int g)^2] = F(\bar{x})^2\mathbb{E}[A^2]/V^2 = F(\bar{x})^2/12$ by Lemma 3.1. Expanding $\big(\int g\,F(u_0 s,\bar{v})\big)^2 = \big(F(\bar{x})\int g + \int g e\big)^2$ and using the two deterministic bounds,

$$\big|\delta_\parallel(p) - \delta_\parallel(0^+)\big| \;\le\; 2\|F\|_\infty\,\mathbb{E}\big[\big|\textstyle\int g\big|\big|\int g e\big|\big] + \mathbb{E}\big[\big(\int g e\big)^2\big] \;\le\; 4\sqrt{2}\,\|F\|_\infty\|\partial F\|_\infty\sqrt{p} + 2\|\partial F\|_\infty^2\,p\,.$$

Since $p = \frac{1}{2}u_0^2 \le L^2/2$ on the diamond, $2\|\partial F\|_\infty^2\,p \le \sqrt{2}\,L\|\partial F\|_\infty^2\sqrt{p}$, which gives the displayed constant. □

Lemma 4.4 (boost-channel annihilation)**.** *Let $\gamma(p;\bar{x})$ be measurable with $|\gamma(p;\bar{x})| \le c\sqrt{p}$ on $(0, p_{\max}(\bar{x})]$, uniformly in $\bar{x}$. Then, under the moment kill $m_2 = 0$,*

$$\rho^2 \int \mathrm{d}\bar{x} \int \mathrm{d}V\,k(\rho V)\,V^2 \int_V^{p_{\max}(\bar{x})} \gamma(p;\bar{x})\,\frac{\mathrm{d}p}{p} \;=\; O\big(\rho^{-3/2}\big),$$

i.e. every boost-channel contribution vanishing at $p = 0$ is annihilated up to relative order $\rho^{-1/2}$; a p-independent contribution $\gamma \equiv \gamma_0(\bar{x})$ instead produces $\int_V^{p_{\max}} \frac{\mathrm{d}p}{p} = \ln\frac{p_{\max}}{V}$, whose V-moment under $m_2 = 0$ is exactly $-m_2^{\log}\,\rho^{-3}\gamma_0$.

Proof. For the first part, extend the inner integral to $(0, p_{\max}]$: the completion error is $\int_0^V |\gamma|\frac{\mathrm{d}p}{p} \le 2c\sqrt{V}$, whose V-integral is an absolute moment $\int V^{5/2}|k(\rho V)|\,\mathrm{d}V = O(\rho^{-7/2})$; the completed integral is a function of $\bar{x}$ alone times $\int V^2 k(\rho V)\mathrm{d}V = m_2\rho^{-3} = 0$. The second part is the computation of Theorem 3.2 in the (V, p) chart. □

Proof of Theorem 4.1. Write $\mathbb{E}[\Phi^2] = V^2\delta_\parallel(p;\bar{x}) + \varepsilon$, with $|\varepsilon| \le 2\sqrt{\mathbb{E}[(\int h_Q F)^2]\,\mathbb{E}[r^2]} + \mathbb{E}[r^2] \le C_F V^2\sqrt{q}$ by Lemma 4.3 and Cauchy–Schwarz. Insert into the Mecke integral in the $(\bar{x}, V, p)$ chart: (i) the q-channel error ε: since $q = V^2/p \le V$ on the chart, $\int_V^{p_{\max}} \sqrt{V^2/p}\,\frac{\mathrm{d}p}{p} \le 2\sqrt{V}$, giving again the absolute $V^{5/2}$-moment, relative order $\rho^{-1/2}$; (ii) the difference channel $\gamma = \delta_\parallel(p) - \delta_\parallel(0^+)$ is $O(\sqrt{p})$ by Lemma 4.3 and is annihilated by Lemma 4.4; (iii) the constant channel $\delta_\parallel(0^+;\bar{x}) = F(\bar{x})^2/12$ produces $-m_2^{\log}/12 \cdot \rho^{-1}\int F^2 = \frac{C_2}{5\rho}\int F^2$; (iv) region-boundary corrections (centers $\bar{x}$ within a ρ-thin corner strip, where $p_{\max}$ degenerates) are controlled exactly as in the constant-F case, contributing the relative $O(\ln(\rho V_D)/(\rho V_D))$. The first-order term needs no separate treatment (Lemma 4.2), which is what permits the C^1 hypothesis. □

Lemma 4.4 is, we believe, the structurally significant statement of this paper: *the coefficient condition that removes the boost-log divergence automatically removes the entire boost-channel dependence on the field profile.* A Lorentz-invariant discrete gauge action is possible because these two requirements — finiteness and boost-uniform locality — are one condition, not two.

5 Variance: the limit holds in L^2

Fluctuations are controlled by the Poincaré inequality for Poisson functionals [26, Thm. 18.7]: $\mathrm{Var}[S] \le \rho \int_D \mathbb{E}[(D_p S)^2]\,\mathrm{d}p$, where $D_p S = S(\eta \cup \{p\}) - S(\eta)$ is the add-one-point cost.

Lemma 5.1 (counting)**.** *Call a pair* (x, y) active *for* p *if* $p \in \{x, y\} \cup I(x, y)$ *and* $n_\eta(I(x, y)\backslash\{p\}) \le 3$. *Then* $|D_p S| \le c\,\Lambda_F^2 \sum_{\text{active}} V_{xy}^2$, *and the expected number of active pairs is* $c_1 \ln(\rho V_D) + O(1)$, *uniformly in* p.

Proof. Adding p changes only: (a) coefficients and chain-pair sums of pairs with $p \in I(x, y)$ — and if $n_\eta \ge 4$ both old and new coefficients vanish (support $\{2, 3\}$), so only $n_\eta \le 3$ contributes, with change bounded by $2 \max |C_n| \cdot \binom{4}{2} \cdot 4\Lambda_F^2 V^2$ (using $|\Phi| \le 2\Lambda_F V$); (b) new pairs with endpoint p, similarly bounded. The count follows from the master corner estimates of Appendix A with $m = 0$: the expected number of sparse intervals through a point is a corner integral with exponentially decaying kernel, whose boost channel contributes one logarithm. □

Lemma 5.2 (add-one cost, second moment)**.** *There is* c *(depending on the* C_n*) such that for* $\rho V_D \ge 2$, *uniformly in* $p \in D$,

$$\mathbb{E}\big[(D_p S)^2\big] \;\le\; c\,\Lambda_F^4\,\frac{\ln^2(\rho V_D)}{\rho^4}\,.$$

The proof — an enumeration of the ways two active pairs can share points, reduced by a geometric-mean bound on joint sparsity and an AM–GM decoupling to products of weighted corner integrals — is elementary but deserves its own space; it is given in Appendix A. This was the one open sub-lemma in an earlier stage of this work; closing it makes the following theorem unconditional.

Theorem 5.3 (variance bound)**.** *For the moment-killed action with* $\Lambda_F = \sup|F|$, $\rho V_D \ge 2$*:*

$$\mathrm{Var}[S] \;\le\; c\,\Lambda_F^4\,V_D\,\frac{\ln^2(\rho V_D)}{\rho^3}, \qquad \textit{hence} \qquad \frac{\mathrm{Var}[S]}{\mathbb{E}[S]^2} \;\le\; C\,\frac{\ln^2(\rho V_D)}{\rho V_D}$$

whenever $\int_D F^2 > 0$ *and* ρ *is large enough that Theorem 4.1 bounds* $\mathbb{E}[S]$ *below.*

Corollary 5.4 (unconditional L^2 limit)**.** *For every* C^1 *field strength on the diamond,*

$$\rho\,S \;\xrightarrow[\rho\to\infty]{L^2}\; \frac{C_2}{5}\int_D F^2\,\mathrm{d}t\,\mathrm{d}x.$$

Remark 5.5 (the $\rho^{-1/3}$ puzzle). Simulations (Section 6) show per-sprinkling fluctuations $\mathrm{std}(r) \propto \rho^{-0.33}$ over a 32× density range — anomalously slow if read as a power law. Theorem 5.3 predicts $\mathrm{std}(r) \lesssim \ln\rho \cdot \rho^{-1/2}$, whose effective log–log slope is $-\frac{1}{2} + 1/\ln\rho \approx -0.34$ at $\rho = 500$ and -0.37 at $\rho = 2000$, so the measured exponent is *consistent with* being the finite-density face of the $\ln\rho \cdot \rho^{-1/2}$ law. Honestly stated, however, a handful of densities over a 32× range cannot discriminate between a naive $\rho^{-1/3}$ power, the predicted $\ln\rho \cdot \rho^{-1/2}$ shape, and nearby alternatives:

ρ	M	r (MC)	r exact (8)	r_{raw} (MC)	r_{raw} exact
250	800	0.1532(120)	0.16702	1.496(13)	1.527
500	500	0.2071(134)	0.18074	2.153(16)	2.132
1000	300	0.1854(145)	0.18898	2.756(16)	2.774
2000	150	0.1756(175)	0.19380	3.446(22)	3.438
4000	72	0.2223(219)	0.19655	4.108(26)	4.115
8000	32	0.1724(200)	0.19810	4.776(27)	4.799

Table 1: Constant-F sweep on the unit diamond ($L = 1$, $E = 10$; M sprinklings per density; seeds fixed and published in the script). The moment-killed column converges to $C_2/5 = 0.2$ along the exact finite-density curve (8); the raw column follows the exact divergent form of Corollary 3.3(iii).

the fluctuation data are consistent with the theorem, not evidence for its sharpness. A further caveat: the constant in Theorem 5.3 is not tracked and is not small — at $\rho = 1000$, $V_D = \frac{1}{2}$, the measured relative fluctuation of a single sprinkling is of order unity while the bound's dimensionless content $\ln^2(\rho V_D)/(\rho V_D) \approx 0.08$, so the implied constant is $\gtrsim 20$, and Corollary 5.4 is an asymptotic statement, empirically untested *as a bound* at simulation-accessible densities. No mesoscale smearing is needed for the convergence statement — though it will be for any fluctuating (quantized) connection, as in the scalar case [4, 9].

6 Numerical confirmation in 1+1 dimensions

All numerics are produced by the ancillary script (`p18_maxwell_causal.py`; fixed published seeds; every quoted number asserted). Sanity checks assert the Stokes identity $(2) = \int_Q F$ to 10^{-12}, gauge invariance of Φ under random gauge functions to 10^{-10}, $\mathbb{E}[A^2] = V^2/12$, and the closed form (6) against direct quadrature of (5) to 5×10^{-7}. Table 1 and Figure 1 present the constant-F sweep: at every density both channels agree with the exact means of Theorem 3.2 within errors (twelve comparisons, pulls between -2.4 and $+2.0$, consistent with unit-normal fluctuations), the raw channel tracking its divergent exact form and the moment-killed channel the convergent one. The precision-weighted mean over $\rho \geq 500$ is $r = 0.1927 \pm 0.0074$, against the identically weighted exact value 0.1894 (asymptote 0.2).

Smooth-field test: for $F(t, x) = E_0[1 + \frac{1}{2}\sin(2\pi t/L)\sin(2\pi x/L)]$ (varying in both coordinates; exact gauge potential in the script), normalized by $\int F^2$: $r = 0.1893 \pm 0.0246$ at $\rho = 1000$ ($M = 130$) and $r = 0.1863 \pm 0.0303$ at $\rho = 4000$ ($M = 36$), consistent with $r \to 0.2$ and with the $O(\rho^{-1/2})$ remainder of Theorem 4.1 — the action localizes on the Maxwell density, not on any diamond average. Fluctuations and the add-one cost are shown in Figure 3: the measured $\mathbb{E}[(D_p S)^2]$ at the diamond center falls by 4.4 orders of magnitude between $\rho = 250$ and $\rho = 4000$ while the normalized ratio $q = \rho^4\,\mathbb{E}[(D_p S)^2]/(E^4 \ln^2(\rho V_D))$ stays flat at 15.5, 14.7, 16.2 — a direct confirmation of the $\ln^2(\rho V_D)\,\rho^{-4}$ law of Lemma 5.2.

An honest caveat carried from the simulations: per-sprinkling action distributions are heavy-tailed (independent batches at the lowest density differed by more than their nominal errors would suggest), so quoted standard errors may be mildly optimistic; the multi-density agreement with the exact curve (8) is the load-bearing check.

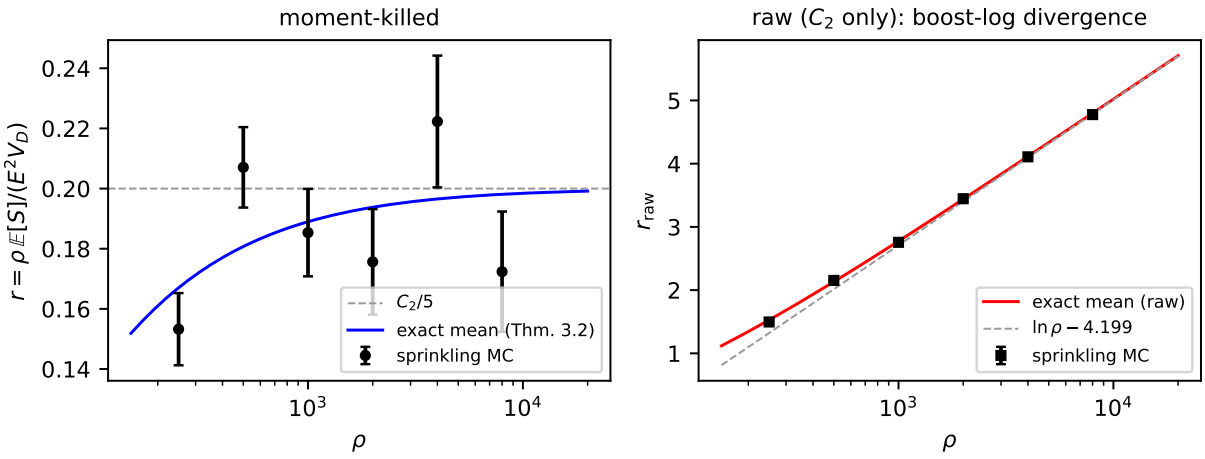


Figure 1: 1+1D convergence. Left: moment-killed normalized mean r versus density; points are sprinkling Monte Carlo, the curve is the exact finite-density mean (8) — every visible finite-size deficit is the computed boundary term. Right: the raw scheme against its exact divergent mean and the asymptote $\ln\rho - 4.199$.

7 2+1 dimensions: derived kill ray and a confirmed coefficient

In $d \geq 3$ the loop bivector $\Sigma^{\mu\nu} = \oint x^{[\mu}\mathrm{d}x^{\nu]}$ is *not* boost invariant: $\mathbb{E}[\Phi^2] = F_{\mu\nu}F_{\rho\sigma}\mathbb{E}[\Sigma^{\mu\nu}\Sigma^{\rho\sigma}]$ grows under boosts of the interval axis at fixed volume, and the boost integral diverges as a power law, which volume moments alone cannot touch — unless the divergent pieces factorize into (boost class)×(volume moment). Two exact structures make this factorization happen.

Proposition 7.1 (endpoint-axis degeneracy)**.** *For a chain-pair loop both endpoints x, y lie on the interval axis. In the rest frame of the interval, every purely spatial component of the loop bivector vanishes identically, $\Sigma^{ij} = 0$ (the spatial projection of the loop is the degenerate path $O \to z_\perp \to O \to z'_\perp$). Consequently the rest-frame correlator has a single constant: in 2+1 dimensions $a_3 := \mathbb{E}[(\Sigma^{xy})^2]/V^{4/3} = 0$ exactly and $b_3 := \mathbb{E}[(\Sigma^{tx})^2]/V^{4/3} = \frac{3}{160}(12/\pi)^{4/3} = 0.111955$ exactly, by the bivector identity of Lemma B.1 below (the ancillary script retains a 4×10^5-sample Monte Carlo cross-check, which agrees at the 0.2% level), and $\mathbb{E}[\Phi^2]$ is a pure frame term in the interval axis.*

The boost covariance $\mathbb{E}[\Phi^2] = V^{4/3}b_3[1+\sinh^2\chi\,\sin^2\varphi]\,E^2$ (axis rapidity χ, azimuth φ; verified by boosted-sample Monte Carlo to 0.3% at $\chi \leq 1.5$), together with the invariant measure $\mathrm{d}y \propto \mathrm{d}V\sinh\chi\,\mathrm{d}\chi\,\mathrm{d}\varphi$ and the exact interval cutoff $\tau_m = A - \sqrt{A^2 - B}$ (whose large-boost expansion, *for the diamond*, contains even powers of $e^{-\chi}$ only, so no fractional moment classes arise; the 3+1-dimensional analogue of this parity structure is proved region-independent across the integer-power class by Lemma 8.4, but its 2+1-dimensional counterpart for general causally convex regions remains an unproved assertion, noted at Open Problem 8.8(c)), organizes the divergences into exactly two classes, $\cosh^3\chi_m$ (volume weight $s^{1/3}$) and $\cosh\chi_m$ (weight s^1); the odd-dimensional boost measure produces no log class. The kill conditions are

$$m_{1/3} = m_1 = 0 \qquad \Longrightarrow \qquad (C_2, C_3, C_4) = (40, -25, 6) \tag{10}$$

(minimal support; integer ray; the BDG-type alternation again fails), and the surviving term gives

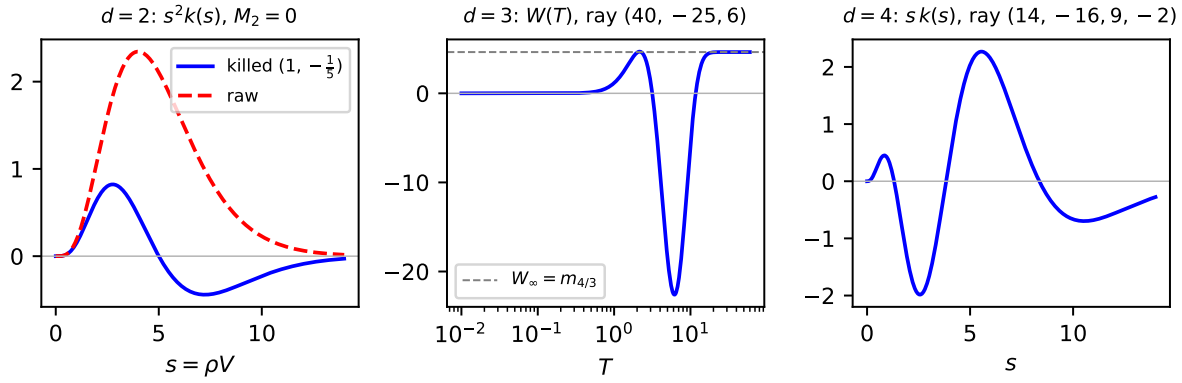


Figure 2: The kill structure in $d = 2, 3, 4$: the weighted kernels whose indicated moments must vanish. Alternation is forced (Corollary 3.3(i) and its analogues): each killed kernel must dip negative.

the prediction

$$r_3 := \rho^{1/3}\,\frac{\mathbb{E}[S]}{E^2 V_D} \xrightarrow[\rho\to\infty]{} \tfrac{8}{3}\,(a_3 - 2b_3)\,m_{4/3} = -\tfrac{16}{3}\,b_3\,m_{4/3} = -2.7647, \tag{11}$$

with $m_{4/3} = \Gamma(13/3)/2 = 4.6303$ for (10) (the ray normalizes the bracket to unity), and convergence $O(\rho^{-1/3})$ — slow, as the one power of $\ell = \rho^{-1/3}$ in $\mathbb{E}[S] \sim \ell\, E^2 V_D$ suggests.

Numerically (script stages `c2p1d`, `exact2p1d`, `sweep2p1d`; Figure 4): the exact reduced mean integral (no sprinkling noise; quasi-Monte Carlo over the region, quadrature over the boost) evaluates to $r_3 = -0.654(17)$, $-0.943(22)$, $-1.221(29)$, $-1.84(4)$, $-1.95(6)$, $-2.26(18)$, $-2.68(46)$, $-2.86(9)$ at $\rho = 1.9 \times 10^3$ through 10^8. The sequence decreases monotonically with measured exponent ≈ 0.3 (predicted $\frac{1}{3}$), and at the highest density sits at $-2.86(9)$ — on the *far* side of the prediction -2.7647, overshooting it by about 1σ. The finite-ρ data are therefore *consistent with the predicted coefficient at the few-percent level but do not yet confirm it*: the residual $O(\rho^{-1/3})$ correction has not been computed, and a one-sided approach cannot be asserted. A structural check, $\int T^{-2} W(T)\,\mathrm{d}T \approx 0$ with W dipping to -22.6 at $T = 6.2$, confirms that killing the $\cosh^3$ class forces the kernel negative. The raw channel diverges as $\rho^{0.98}$ and is reproduced by sprinkling Monte Carlo at the one-to-two-percent level ($2.601(14)$, $5.011(25)$, $9.48(8)$ against exact 2.559, 4.979, 9.543 at $N = 500, 1000, 2000$), validating measure, covariance and cutoff against real sprinklings — but *not* the kill conditions, since nothing is killed in the raw channel; the killed sprinkling channel is noisy (the alternating weights amplify fluctuations) and consistent within its large errors: -0.31 ± 0.78, -1.7 ± 1.1, -4.5 ± 5.7 against exact -0.654, -0.943, -1.221. Honest scope: this is a pilot — mean only, constant field, one region, with the smooth-field and variance analyses of Sections 4–5 expected but not proven to lift.

8 3+1 dimensions: the constant-field Maxwell theorem

The 3+1-dimensional mean is now a theorem for constant field strength on a causal diamond — the region qualifier is load-bearing and is discussed after the theorem. Proposition 7.1 holds verbatim and sharpens: all spatial–spatial components vanish *per loop*, and the surviving components are exactly linear in the intermediates (Lemma B.1), giving the rest-frame constant $b_4 = \mathbb{E}[(\Sigma^{0i})^2]/V =$

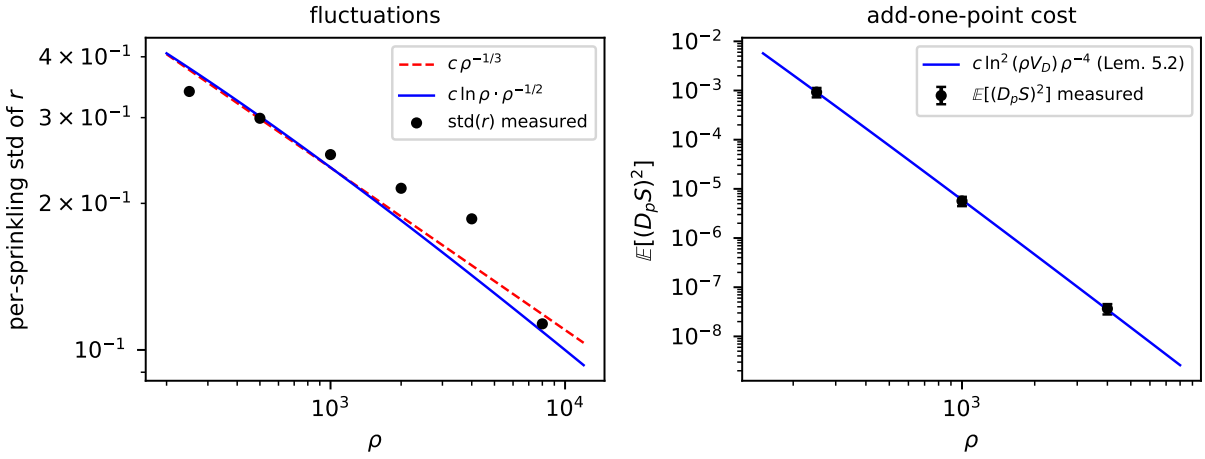


Figure 3: Left: per-sprinkling fluctuation of r; the $\ln\rho\cdot\rho^{-1/2}$ law of Theorem 5.3 and the naive $\rho^{-1/3}$ power fit are indistinguishable over this range (Remark 5.5). Right: second moment of the add-one-point cost at the diamond center versus the $\ln^2(\rho V_D)\rho^{-4}$ law of Lemma 5.2 (curve normalized to the data mean).

$2/(5\pi) = 0.127324$ in closed form (the Monte Carlo value 0.1277 of an earlier stage of this work was this number), so

$$\mathbb{E}[\Phi^2] = b_4\, V\, E_n^2, \qquad E_n^2 = (F\eta F)_{\alpha\beta}\, n^\alpha n^\beta$$

— the squared electric field in the frame of the interval axis n: a pure frame term. The feared mixing of the two quartic invariants does not arise for chain-pair loops; in the basis $\{F_{\mu\nu}F^{\mu\nu}, (F\cdot n)^2, (\tilde{F}\cdot n)^2\}$ the correlator has coefficients exactly $(0, *, 0)$ — the magnetic structure decouples identically. The Poisson pair ensemble then averages this frame term over interval axes with the invariant hyperboloid measure $\mathrm{d}V\sinh^2\chi\,\mathrm{d}\chi\,\mathrm{d}\Omega_2$ cut off by the region. Crucially, in the proof of Theorem 8.2 the axis average is performed *after* the region integral, whose covariogram is independent of the spatial axis direction: the average that actually occurs is a compact, exact S^2 integral with the rapidity range cut off by the region itself, and no regulator or non-compact Lorentz average appears anywhere (Appendix B, Step 3).

Remark 8.1 (heuristic isotropy). The following recapitulates, at the level of a regulated axis average, the isotropy that the exact region integral of Appendix B establishes on its own. It is *not* load-bearing — no regulated average occurs in the proof of Theorem 8.2 — and is retained only because it exhibits the cancellation mechanism in closed form. Write the regulated axis average as $\int_{n^0\le X} n^\alpha n^\beta\,\mathrm{d}\Omega_H = \mathcal{A}(X)\,\eta^{\alpha\beta} + \mathcal{B}(X)\,\hat{t}^\alpha\hat{t}^\beta$. The power-divergent parts of $\mathcal{A}, \mathcal{B}$ are frame dependent but lie in killed moment classes; the χ-linear (logarithmic) parts satisfy

$$\mathcal{A}_{\log} = -\tfrac{\pi}{2}\,\chi, \qquad \mathcal{B}_{\log} = 0 \quad \text{identically}$$

(from $\int\sinh^2\chi\,\mathrm{d}\chi \supset -\frac{\chi}{2}$ and $\int\sinh^2\chi\cosh^2\chi\,\mathrm{d}\chi \supset -\frac{\chi}{8}$: $4\mathcal{A}+\mathcal{B} = -2\pi\chi$, $\mathcal{A}+\mathcal{B} = -\frac{\pi}{2}\chi$, whence $\mathcal{B}_{\log} = 0$). Since only the log class survives the kills below, and the log coefficient of the cutoff is direction independent, the surviving mean is exactly η-contracted: $(Fn)(Fn) \to \propto F_{\mu\nu}F^{\mu\nu}$ — at this heuristic level no frame residue survives. The theorem-level statement of the same fact is the exact η-contraction obtained in Appendix B, Step 3.

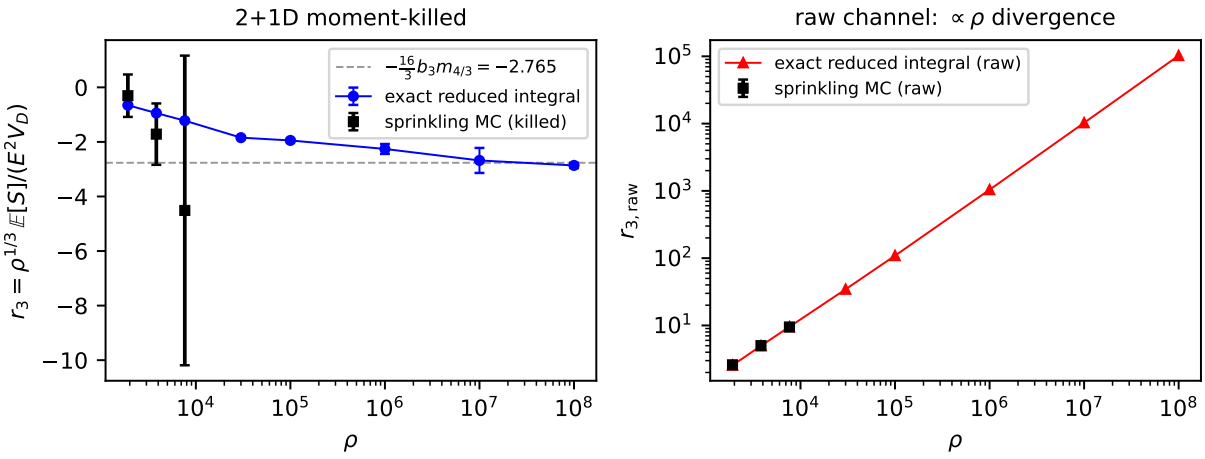


Figure 4: 2+1D pilot. Left: killed channel — exact reduced integral (circles) approaching the predicted coefficient -2.765 (dashed), with sprinkling Monte Carlo (squares). Right: raw channel — the $\propto \rho$ divergence, exact integral versus sprinkling Monte Carlo.

The divergence classes are $\cosh^4 \chi_m$ (s^0), $\cosh^2 \chi_m$ ($s^{1/2}$) and the χ-linear class (s^1, log moment); the kill system and its minimal integer solution are

$$m_0 = m_{1/2} = m_1 = 0: \qquad (C_2, C_3, C_4, C_5) = (14, -16, 9, -2), \qquad m_1^{\log} = -\tfrac{1}{2} \text{ exactly} \tag{12}$$

(the rational value follows from $m_1 = 0$, which cancels the Euler-constant terms, and $42H_3 - 192H_4 + 270H_5 - 120H_6 = -\frac{1}{2}$; asserted symbolically in the script). Dimensional bonus: in $d = 4$ the mean action is *scale-free*, $\rho^2 \cdot \rho^{-2} = \rho^0$: $\mathbb{E}[S] \to c \int F_{\mu\nu} F^{\mu\nu}\, \mathrm{d}^4x$ with c a pure number $\propto b_4\, |m_1^{\log}|$ — Maxwell marginality emerges from causal counting. One coincidence should be pre-empted: the exact $d = 4$ BDG layer coefficients are $(1, -9, 16, -8)$ [5, 7], and the ray (12) is $(14, -16, 9, -2)$; the shared digits $\{9, 16\}$ are coincidental. The two vectors solve different linear systems (BDG: layer-moment cancellation so the sprinkling mean recovers $\Box\phi$; here: $\binom{n}{2}$-weighted Gamma-moment cancellation for the pair-count kernel (4)); the BDG vector fails the present kill system in every placement (on support $n = 2, \ldots, 5$ it gives $m_0 = -10$, $m_1 = -105$), and no permutation-plus-affine map carries one vector to the other (checked symbolically over all 24 permutations).

An earlier stage of this work left the assembly of these structures as an open problem, conjecturing a finite coefficient $c = \kappa\, b_4 |m_1^{\log}|$ and a boundary estimate of polylog$/(\rho V_D)^{1/4}$ class. Both parts are now closed *for the causal diamond* — the second in a stronger form than conjectured, because a covariogram identity special to the causal diamond — the Lorentzian analogue of Matheron's set covariogram from Euclidean integral geometry [22, 23], elementary here because light-cone nesting collapses two of the four constraints — makes the region integral *exact* (Lemma B.2): the mean collapses to a closed-form double integral, the 4-dimensional analogue of Theorem 3.2.

Theorem 8.2 (3+1D Maxwell mean on a causal diamond, constant F)**.** *Let D be the causal diamond of height T in $\mathbb{M}_4$ ($V_D = \pi T^4/24$), F a constant field strength, and S the action* (3) *with the killed ray* (12)*. Then, up to $O(e^{-c\rho V_D}\,\mathrm{poly})$ moment-truncation tails, $\mathbb{E}[S]$ is given exactly at every density by the finite closed form* (25)–(26)*; the only divergence classes of the general-coefficient mean are $\rho\, m_0$, $\rho^{1/2} m_{1/2}$ and $\ln \rho \cdot m_1$ (quarter-power classes cancel identically), so the kill system* (12) *is sufficient for every constant F, and necessary for every F outside the measure-zero set $\{4B^2 = E^2\}$, on which the $\rho^{1/2}$ class is annihilated by the tensor structure rather than by*

the coefficients (the m_0 condition is unconditionally necessary, since its tensor factor is $M_{00}+\bar{M} = \frac{2}{3}(E^2+B^2) > 0$ for $F \neq 0$); and

$$\mathbb{E}[S] \xrightarrow[\rho\to\infty]{} \frac{3}{20\pi}\int_D F_{\mu\nu}F^{\mu\nu}\,\mathrm{d}^4x, \qquad \frac{3}{20\pi} = \kappa\, b_4\,\big|m_1^{\log}\big|, \quad \kappa = \tfrac{3}{4},\ b_4 = \tfrac{2}{5\pi},\ m_1^{\log} = -\tfrac{1}{2},$$

for every constant F: the surviving mean is exactly η-contracted (no frame residue), and the approach is

$$r_4(\rho) := \frac{\mathbb{E}[S]}{E^2 V_D} = -\frac{3}{10\pi} + \frac{\sqrt{6}}{32\pi}\,\frac{105\big[\ln(\pi\rho/6)+\gamma_E\big]-737}{\sqrt{\rho}} + O\Big(\frac{\ln\rho}{\rho}\Big) \tag{13}$$

for the pure-electric normalization ($T = 1$; $F_{\mu\nu}F^{\mu\nu} = -2E^2$).

The proof — the bivector identity, the covariogram lemma, the exact double integral, its evaluation by an exact Laurent-plus-log integral and a moment mapping, and the class table — is given in Appendix B, together with a four-tier verification hierarchy (Step 4) whose primary numerical evidence is an end-to-end Mecke Monte Carlo using none of the appendix's lemmas.

Scope: which regions. Theorem 8.2 is proved on the causal diamond, and its exact finite-density closed form (25)–(26) is diamond-specific: the covariogram polynomial $\frac{\pi}{24}(T-u)^2(T-v)^2$ belongs to that one region. The class hypothesis of Lemma 8.4 below — a single absolutely convergent power series on a full box, with the datum vanishing outside — is satisfied by the causal diamond and essentially by it alone: it requires both a box-shaped (u,v) support (equivalently, a boundary whose maximising locus slides along a null generator in every spatial direction) and the absence of interior critical values of the covariogram's level function; generic regions fail both (Proposition 8.6(i)). That hypothesis is, however, stronger than the argument uses. The boost integral at fixed proper time is a function of $uv = \omega^2$ alone for every compact region — but that alone does not force even powers of ω ($|\omega|$ is a function of ω^2); what does is integer-power expandability of the datum in the null coordinates, and granted it, the $u \leftrightarrow v$ parity, the odd-power structure and the absence of a quarter-power class follow with no further class hypothesis, and the ρ^0 log coefficient is fixed by the degree-2 jet, hence by $\mathrm{Vol}(R)$ alone (Proposition 8.6(ii),(iii)). That expandability is exactly what the argument requires, as regularity at *null* displacements: that $g_R(u,v)$ admit an integer-power expansion to second order in v as $v \to 0^+$ for almost every u in the null range — condition (M). That condition holds exactly for the causal diamond, for every region with piecewise-linear null-or-spacelike boundary (domains of dependence of polytopes, truncated diamonds, finite intersections of diamonds), and for the exactly solvable smooth family, and is confirmed numerically on a six-region library; the half-integer and logarithmic loci that generic smooth regions do carry — a $(\Lambda - w^0)^{(n+2)/2}$ edge at the support boundary and $|w^0-\lambda_c|^{(n+2)/2}$ interior Morse loci — are integrated over by the boost measure and affect coefficients only. Two structural facts should also be recorded (Lemma 8.5): a compact causally convex region has achronal boundary, so it is exactly a region $\{f^- \le t \le f^+\}$ with $f^\pm$ 1-Lipschitz — in particular solid cylinders, tubes and slabs are never causally convex, and are not hard cases but non-cases — and its covariogram is $g_R(w) = \int_{K\cap(K-\vec{w})}[f^+(x+\vec{w}) - f^-(x) - w^0]_+$, whose first-order term is an exact linear form in (u,v): the Lorentzian analogue of Matheron's perimeter term, which in Euclidean signature is the non-analytic $c_1|w|$. We therefore state the result as: *the action defined by the ray $(14,-16,9,-2)$ has a Maxwell mean, with the same constant $\frac{3}{20\pi}\mathrm{Vol}(R)$, on every compact causally convex region satisfying the null-displacement condition (M) — verified for the causal diamond, for the polytopal family and for the smooth family; the exact finite-density closed form remains diamond-specific.* The analogous statement in 2+1 dimensions is untreated, and there the Morse classification predicts logarithmic rather than half-integer interior loci, which is worth

checking before the 2+1-dimensional ray $(40, -25, 6)$ is called region-independent. A related honest caveat stands: the coefficient ray is dimension-dependent ($(1, -\frac{1}{5})$ in 1+1D, $(40, -25, 6)$ in 2+1D, $(14, -16, 9, -2)$ in 3+1D), so the construction supplies a family of actions re-derived per dimension — though, within the class (M), no longer per region.

Remark 8.3 (no quarter-power boundary class). The finite-size corrections in (13) carry $\ln\rho\cdot\rho^{-1/2}$, not the conjectured $(\rho V_D)^{-1/4}$: the quarter-moment classes cancel identically in the exact region integral of the diamond — an instance of the parity mechanism of Lemma 8.4 below. Numerically the approach mimics an effective $\rho^{-1/4}$ power over the sprinkling-accessible range — the 3+1-dimensional reprise of the $\rho^{-1/3}$ puzzle of Remark 5.5.

The quarter-power cancellation is not an accident of the diamond. The following lemma — prompted by an observation of an anonymous adversarial reviewer, and verified symbolically by the ancillary script (stage `parity`) — isolates the mechanism and shows that both the kill system and the continuum constant are independent of the region across an explicitly characterized class.

Lemma 8.4 (parity; region-independence of the kill system and the constant)**.** *Let $R \subset \mathbb{M}_4$ be compact and causally convex with piecewise-smooth boundary, $g_R(w) = \mathrm{Vol}(R \cap (R - w))$ its covariogram, and for future-timelike $w = \tau n$ (null coordinates $u = \tau e^{\chi}$, $v = \tau e^{-\chi}$, axis $n = (\cosh\chi, \sinh\chi\,\hat{\omega})$, $\chi > 0$) let*

$$H_R(u,v) \;=\; \int_{S^2} g_R(w)\, M_{\alpha\beta}\, n^{\alpha} n^{\beta}\, \mathrm{d}\Omega_2, \qquad M_{\alpha\beta} = (F\eta F)_{\alpha\beta} \text{ (notation of Appendix B)},$$

so that, by the Mecke reduction, Lemma B.1 and $\mathrm{d}^4 w = \frac{1}{8}(u-v)^2\,\mathrm{d}u\,\mathrm{d}v\,\mathrm{d}\Omega_2$,

$$\mathbb{E}[S] \;=\; \frac{\rho^2 b_4}{16} \int_0^{\infty}\!\!\int_0^{\infty} \mathrm{d}u\,\mathrm{d}v\; (u-v)^2\, V\, k(\rho V)\, H_R(u,v), \qquad V = \frac{\pi u^2 v^2}{24} \tag{14}$$

(symmetrized to the full quadrant by part (i)). Say R is of integer-power class *if there are $T_R > 0$ and symmetric coefficients $p_{ab} = p_{ba}$ with $\sum_{a,b} |p_{ab}| T_R^{a+b} < \infty$ such that*

$$\frac{uv}{\pi}\, H_R(u,v) = \sum_{a,b\geq 0} p_{ab}\, u^a v^b \quad \text{for } \max(u,v) < T_R, \qquad H_R = 0 \text{ for } \max(u,v) \geq T_R.$$

Then:

(i) The $u \leftrightarrow v$ symmetry is free*: $H_R(u,v) = H_R(v,u)$ for* every *compact R — no rotation-invariance hypothesis enters. ($u \leftrightarrow v$ is $\vec{w} \mapsto -\vec{w}$ at fixed w^0, which the S^2 integral absorbs by $\hat{\omega} \mapsto -\hat{\omega}$.)*

(ii) Only odd powers; the same kill system. *For R of integer-power class, the substitution $u = \omega\mu$, $v = \omega/\mu$ turns* (14) *into exact μ-integrals over the $\mu \mapsto 1/\mu$-symmetric range $(\omega/T_R, T_R/\omega)$; a monomial $u^a v^b$ of the expanded integrand survives only at the* odd *powers ω^{2b+1} and ω^{2a+1}, with a logarithm exactly on the diagonal $a = b$. Under the moment map of Appendix B, Step 3, the divergence classes are therefore exactly $\rho\, m_0$, $\rho^{1/2} m_{1/2}$ and $\ln\rho\cdot m_1$ — no quarter-power class arises for any region of the class — and the kill system* (12) *with the ray $(14, -16, 9, -2)$ suffices for every such R.*

(iii) The constant depends on the region only through its volume. *For every compact R, the degree-2 jet of uvH_R/π is $\mathrm{Vol}(R)\big[M_{00}(u+v)^2 + \bar{M}(u-v)^2\big]$; the ρ^0 log class sits on the*

diagonal $a = b = 3$ of the degree-6 core, which only this jet reaches, with diagonal coefficient $(6\bar{M} - 2M_{00})\,\mathrm{Vol}(R)$. Hence, for every R of integer-power class, with the killed ray,

$$\mathbb{E}[S] \;\xrightarrow[\rho\to\infty]{}\; \frac{3\,(3\bar{M} - M_{00})}{20\pi}\,\mathrm{Vol}(R) \;=\; \frac{3}{20\pi}\int_R F_{\mu\nu}F^{\mu\nu}\,\mathrm{d}^4x$$

for every constant F: the limit of Theorem 8.2 — though not its exact finite-density closed form — is region-independent on the class.

Proof. (i) Fix (u, v), i.e. (τ, χ). Since $n(-\chi, \hat{\omega}) = n(\chi, -\hat{\omega})$, exchanging $u \leftrightarrow v$ replaces the integrand of H_R by its value at $\hat{\omega} \mapsto -\hat{\omega}$, and the S^2 measure is reflection invariant.

(ii) The support condition is exactly the μ-window: at fixed ω, $u < T_R \iff \mu < T_R/\omega$ and $v < T_R \iff \mu > \omega/T_R$. With $\mathrm{d}u\,\mathrm{d}v = 2\omega\,\mathrm{d}\omega\,\mathrm{d}\mu/\mu$, the integrand of (14) is $\rho^2 b_4 \frac{\pi^2}{384}\,uv(u-v)^2\,k \times (uvH_R/\pi)$; expanding $uv(u-v)^2\sum p_{ab}u^a v^b$ into monomials $u^{a'}v^{b'}$ and substituting, each monomial contributes $\omega^{a'+b'+1}\mu^{a'-b'}$, and

$$\int_{\omega/T_R}^{T_R/\omega} \mu^{\,n}\,\frac{\mathrm{d}\mu}{\mu} = \frac{(T_R/\omega)^n - (\omega/T_R)^n}{n} \quad (n = a' - b' \neq 0), \qquad = \; 2\ln\frac{T_R}{\omega} \quad (n = 0),$$

so the surviving powers are $\omega^{a'+b'+1\mp n} = \omega^{2b'+1},\ \omega^{2a'+1}$ — odd for every monomial — with the logarithm exactly at $a' = b'$ (absolute convergence on $\max(u, v) \le T_R$ justifies term-by-term integration). The moment map $\omega^g \mapsto \rho^{2-(g+1)/4}\,m_{(g-3)/4}$ then meets only $g = 3, 5, 7$ as non-convergent classes: $\rho\,m_0$, $\rho^{1/2}m_{1/2}$, and the ρ^0 class carrying m_1 and $m_1^{\log}$.

(iii) As $w \to 0$, $g_R(w) = \mathrm{Vol}(R) + O(|w|)$ (the covariogram of a compact set with piecewise-smooth boundary is Lipschitz), and $\int_{S^2} Mnn\,\mathrm{d}\Omega_2 = 4\pi[M_{00}\cosh^2\chi + \bar{M}\sinh^2\chi]$ exactly, so

$$\frac{uv}{\pi}H_R(u, v) = \mathrm{Vol}(R)\big[M_{00}(u+v)^2 + \bar{M}(u-v)^2\big] + O\big((u+v)^3\big)\,:$$

the degree-2 coefficients of the expansion are the displayed jet. The core $uv(u-v)^2\big[M_{00}(u+v)^2 + \bar{M}(u-v)^2\big]$ is homogeneous of degree 6 with diagonal $(3,3)$ coefficient $-2M_{00} + 6\bar{M}$; every higher-degree term of the expansion lands at $\omega^{\ge 9}$, i.e. in convergent classes. Repeating the Step-3 assembly of Appendix B verbatim with the diamond's $c_{00} = T^4$ replaced by $(24/\pi)\mathrm{Vol}(R)$ gives, on the killed ray with $m_1^{\log} = -\frac{1}{2}$, the displayed limit; for the diamond $\mathrm{Vol}(D) = \pi T^4/24$ recovers Theorem 8.2. The stage `parity` of the ancillary script verifies symbolically the μ-integral parity, the log coefficients $(8c_{00}(M_{00} - 3\bar{M})$ at ω^7 and $8c_{11}(M_{00} - 3\bar{M})$ at ω^9 for a generic symmetric factor $c_{00} + c_{10}(u+v) + c_{11}uv$; all other log coefficients zero), the diamond class table, and the assembled region-independent constant. □

Lemma 8.4's hypothesis is, however, stronger than its proof needs, and the gap decides how far Theorem 8.2 reaches. Two structural facts about compact causally convex regions come first; they are elementary, but we have not found them stated together, and each is load-bearing below.

Lemma 8.5 (structure and covariogram of compact causally convex regions)**.** *Let $R \subset \mathbb{M}_{1+n}$ be compact with nonempty interior and $g_R(w) = \mathrm{Vol}(R \cap (R - w))$.*

(i) Structure. *R is causally convex if and only if $R = \{(t, x) : f^-(x) \le t \le f^+(x)\}$ with $f^\pm$ 1-Lipschitz on the compact spatial shadow $K = \{f^- \le f^+\}$. Hence ∂R is achronal, every boundary normal is causal, and no compact causally convex region has a timelike boundary portion. In particular solid cylinders, tubes and slabs are* never *causally convex, and a round ball $\{t^2 + |x|^2 \le R^2\}$ is not either (there $|\nabla f| > 1$ for $|x| > R/\sqrt{2}$; its causal completion is the diamond). The (u, v) support is downward closed and $u \leftrightarrow v$ symmetric, so the boost window at fixed $\hat{\omega}$ is always a single $\mu \mapsto 1/\mu$-symmetric interval.*

(ii) Master formula. *For every future-timelike w,*

$$g_R(w) = \int_{K\cap(K-\vec{w})} \big[f^+(x+\vec{w}) - f^-(x) - w^0\big]_+ \,\mathrm{d}^n x, \tag{15}$$

the fibre's min/max switches being inactive because $f^\pm$ are 1-Lipschitz and $|\vec{w}| < w^0$. For the causal diamond (15) *returns $\frac{\pi}{24}(T-u)^2(T-v)^2$, whose monomials $u^a v^b$ have $a, b \in \{0,1,2\}$: total degrees $0,1,2,3,4$, all integer,* including *a degree-1 term.*

(iii) The Lorentzian perimeter term. *From* (15), $-\partial_{w^0} g_R|_0 = |K|$ *and* $\nabla_{\vec{w}} g_R|_0 = \int_K \nabla f^+ =: \vec{c}$, *so for every compact causally convex R*

$$g_R = \mathrm{Vol}(R) - \tfrac{1}{2}\big[|K|(u+v) - (\vec{c}\cdot\hat{\omega})(u-v)\big] + O(\textit{degree } 2).$$

The first-order term is an exact linear form *in (u,v) — an integer power — and this is special to Lorentzian signature. Matheron's Euclidean perimeter term $-\frac{1}{2}\int_{\partial R} |\langle \nu, w\rangle| \,\mathrm{d}A$ is genuinely non-analytic (it is $c_1|w|$ for a ball) because the normals sweep a continuum and the kink locus fills space. Here ν is causal by (i) and w is timelike, so $\langle \nu, w\rangle = \frac{1}{2}[(\nu^0 + \vec{\nu}\cdot\hat{\omega})\,u + (\nu^0 - \vec{\nu}\cdot\hat{\omega})\,v]$ with* both *coefficients carrying the sign of ν^0: the absolute value never kinks on the open quadrant $u, v > 0$, the kink locus being pushed onto the null boundary $u = 0$ or $v = 0$. Causal convexity is precisely what analytises Matheron's perimeter term.*

Proof. (i) Vertical segments are timelike, so every fibre is an interval $[f^-(x), f^+(x)]$. If $f^+(y) - f^+(x) > |y - x|$ then, with $p = (f^+(x), x)$ and $q = (f^+(y), y)$, $q \in I^+(p)$ and $I(p,q)$ contains points just above f^+ near x, outside R. Conversely, for 1-Lipschitz $f^\pm$ and $r = (t,z) \in I(p,q)$ with $p, q \in R$, $t \le t_q - |z-y| \le f^+(y) - |z-y| \le f^+(z)$, and symmetrically below. The named counterexamples are the statement applied to a timelike wall. Downward closure of the support is causal convexity again: if $x, x+w \in R$ then $I(x, x+w) \subset R$, so $x + w' \in R$ for $0 \le w' \le w$. (ii) The fibre of $R \cap (R - w)$ over x is $[\max(f^-(x), f^-(x+\vec{w}) - w^0), \min(f^+(x), f^+(x+\vec{w}) - w^0)]$; the Lipschitz bound gives $f^\pm(x+\vec{w}) - f^\pm(x) \le |\vec{w}| < w^0$, which switches both extrema off. For the diamond $f^\pm = \pm(\rho - |x|)$ in $n = 3$, the superlevel sets of $\Theta = 2\rho - |x| - |x+\vec{w}|$ are prolate spheroids and the λ-integral of $(4\pi/3)a(a^2 - s^2)$ evaluates to the stated polynomial. (iii) Differentiate (15) at $w = 0$, where the positive part is active exactly on $\{f^- < f^+\}$; the signature statement is the displayed identity for $\langle \nu, w\rangle$. Stage `region` of the fourth ancillary script verifies (i)'s cylinder counterexample, (ii) symbolically for the diamond, and (iii) against $|K|$ on a six-region library. □

Proposition 8.6 (the region class, characterized)**.** *With the notation of Lemma 8.4:*

(i) The literal class is the diamond. *Membership in the integer-power class as defined — one absolutely convergent power series on a full box, the datum vanishing outside — forces two independent non-generic properties: a box-shaped (u,v) support (equivalently $\Lambda'(S) \equiv -1$ in every spatial direction, i.e. the maximising locus of $\Theta_{\vec{w}}$ slides along a null generator), and no interior critical values of $\Theta_{\vec{w}}$. The diamond has both — it is the degenerate case, $|x| + |x+\vec{w}|$ attaining its minimum on the whole segment $[-\vec{w}, 0]$. Polytopal regions fail the first (their supports are triangles, or piecewise-linear), smooth regions fail both. So Open Problem 8.8(c)* as posed *has the answer* no. *(This and the region-independence headline of Lemma 8.4 describe the same narrow class from opposite ends: the literal integer-power class is essentially the diamond alone, and the region-independence content survives because the operative condition is (M) of part (iv), not the class hypothesis.)*

(ii) Parity, given (M). *Writing $u = \omega\mu$ and $v = \omega/\mu$, the boost datum is*

$$\Psi_R(\omega) = \iint F_R(u,v)\,\delta(uv-\omega^2)\,\mathrm{d}u\,\mathrm{d}v,$$

a function of ω^2 alone — for every *compact R, whatever its support shape and however non-analytic g_R is. That alone does* not *force even powers: $|\omega| = \sqrt{\omega^2}$ is a function of ω^2, and if F_R carried $(uv)^{1/2}\times$(integer-power smooth) then $\Psi_R \propto \omega\times$(even) and a $\rho^{3/4}$ class would arise — exactly the structure the numerical control below breaks at ($\sqrt{uv}\times$smooth, relative residual 2.3×10^{-3}). What forces even powers is integer-power expandability of F_R in the null coordinates, i.e. condition (M) of part (iv). Granted (M), Ψ_R carries only even powers, $G(\omega) \propto \omega^3\Psi_R$ only odd powers, the divergence classes are exactly $\rho\, m_0$ (ω^3), $\rho^{1/2}m_{1/2}$ (ω^5) and the ρ^0 class (ω^7), and no fractional-power class can arise unless F_R itself carries a* non-integer *power of uv — a half power of uv gives the $\rho^{3/4}$ class above; a quarter power gives Puiseux exponents with denominator 4. No class hypothesis and no rotation hypothesis enter beyond (M).*

(iii) The constant, given (M). *The ρ^0 log coefficient is set by the u^2v^2 coefficient of F_R; since the prefactor $(u-v)^2[M_{00}(u+v)^2+\bar{M}(u-v)^2]$ is exactly degree-4 homogeneous, that coefficient is prefactor$\times g_R(0,0) =$ prefactor$\times\mathrm{Vol}(R)$ and nothing else. Lemma 8.4(iii) therefore holds with (M) in place of the integer-power-class hypothesis — conditional on (M), not unconditional. (Numerically: the $\omega^4\ln(1/\omega)$ coefficient is $-4\,g_R(0,0)$ to ten digits for triangular support, box support, analytic data, half-integer-power data and data with a nonzero uv jet, and doubles exactly when $g_R(0,0)$ does.)*

(iv) The operative condition is regularity at null displacements. *The small-ω expansion of Ψ_R is governed by the poles at $z = 0,-1,-2$ of the Mellin transform $\iint F_R(u,v)(uv)^{z-1}\mathrm{d}u\,\mathrm{d}v$, which are located solely by the behaviour of F_R as $v\to0^+$ at fixed u (and symmetrically). Hence what the argument needs is*

> ***(M)*** for a.e. $\hat{\omega}$, $g_R(u,v,\hat{\omega})$ admits an expansion in *integer* powers of v to order v^2 as $v\to0^+$, for a.e. u in the null range, with coefficients integrable against the u-weights (and symmetrically in $u\leftrightarrow v$).

Non-analyticity in the interior *of the support, or on the support boundary, is integrated over by the boost measure and affects coefficients only. Condition (M) holds: exactly for the causal diamond (g_D polynomial); for every region with piecewise-linear null-or-spacelike boundary — g_R is then piecewise polynomial with straight break lines, so one-sided integer v-expansions at $v=0^+$ always exist; and for the exactly solvable smooth family $f^{\pm} = \pm(\rho^2-|x|^2)/2\rho$, whose covariogram is $(8\pi/15\rho)B^{5/2}$ with B affine in v at fixed $u<u_*$. Under (M) the conclusions of Lemma 8.4(ii),(iii) hold verbatim.*

Proof. (i) is the definition unwound through (15): by the layer-cake form, the non-analytic loci of g_R in the (u,v) quadrant are exactly the critical values of $\Theta_{\vec{w}}$, which in n spatial dimensions contribute $(\Lambda-w^0)^{(n+2)/2}$ at the support boundary (global maximum) and $|w^0-\lambda_c|^{(n+2)/2}$ at interior nondegenerate critical values (n odd; logarithms for n even); in 3+1 dimensions both exponents are $\frac{5}{2}$. Box support requires support-boundary slope exactly -1; absence of interior critical values requires a degenerate critical manifold. (ii) the displayed δ-representation is immediate from $\mathrm{d}u\,\mathrm{d}v = 2\omega\,\mathrm{d}\omega\,\mathrm{d}\mu/\mu$ and needs no hypothesis on F_R beyond integrability; the passage from function-of-ω^2 to even powers is where (M) enters, through the Mellin pole location of (iv) —

integer-power behaviour at $v \to 0^+$ places the poles at integer z and forbids the ω^{2k+1} terms. (iii) follows from (ii) plus exact degree-4 homogeneity of the prefactor, since only the u^2v^2 coefficient can reach ω^4 on the diagonal. (iv) The Mellin location statement is standard; the three families are the computations indicated. All four parts, together with the decisive numerical discriminator — an even-only-plus-log basis against one augmented with $\{\omega, \omega^3, \omega^{1/2}, \omega^{3/2}\}$, fitted at 50 digits on 16 geometrically spaced ω — are run by the stage `region` of the fourth ancillary script: relative residuals $\leq 10^{-20}$ and fitted odd/half coefficients $\leq 10^{-13}$ for polynomial data on triangular and box support, for support-boundary exponents $\frac{3}{2}, \frac{5}{2}, \frac{7}{2}$, for an interior Morse locus $|u+v-\lambda_c|^{5/2}$ and for a polytopal interior kink, against a control $\sqrt{uv}\times$smooth — a violation of (M) — which breaks the structure at relative residual 2.3×10^{-3}. Only a non-integer power of the null coordinates themselves breaks parity. □

Remark 8.7 (null boundaries do not break the class). The delicate loci named in earlier versions of this paper — null boundary segments and tips — turn out not to be delicate. The diamond's boundary is *entirely* null and its covariogram is a polynomial; flat null faces give support-boundary slope -1 and linear sliding, again integer. What the boost integral is sensitive to is regularity at null *displacements*, condition (M), which is a different statement, and it holds on every family examined. The half-integer and logarithmic non-analyticities that real regions do carry sit at the support boundary and at interior Morse critical values, and *are integrated over*: fed through the boost measure they return integer powers of ω. For the support-edge singularity this is proved, not merely measured: with $F = (c-u-v)^p$ one has exactly $\Psi_p(\omega)/2 = (c-2\omega)^p \int_0^1 (1-\sigma)^p \sigma^{-1/2}(\sigma+\delta)^{-1/2} \mathrm{d}\sigma$ with $\delta = 4\omega/(c-2\omega)$, and the σ-integral's $\delta \ln(1/\delta)$ coefficient is exactly $p/2$ — precisely the value that cancels the odd-power logarithms generated by $(c-2\omega)^p$, leaving $\Psi_p = A(\omega) + B(\omega)\ln(1/\omega)$ with A, B analytic and even.

What remains open in 3+1 dimensions is the smooth-field extension and the residue of the membership question, which we pose precisely:

Open Problem 8.8 (3+1D smooth fields; region independence)**.** *Extend Theorem 8.2 to C^1 field strengths, i.e. prove $\mathbb{E}[S] \to \frac{3}{20\pi}\int_D F_{\mu\nu}F^{\mu\nu}$ for varying F. The central-symmetry kill of the first-order Taylor term (Lemma 4.2) lifts verbatim — it is measure-theoretic, and in 4D the point reflection alone sends $\Sigma \to -\Sigma$ while preserving the ∂F-flux term, so the first-order cross term again averages to zero (as in 1+1D, composition with $z \leftrightarrow z'$ is not needed). What is missing is (a) the 4D analogue of the null-strip decomposition (Lemmas 4.3–4.4): show that field-profile contributions along the stretched null direction enter the exact (u, v) pair measure of Appendix B only with the v-weights v^1, v^2, v^3 (killed by $m_0, m_{1/2}, m_1$ respectively) plus $O(\sqrt{v})$-suppressed remainders; (b) an Item-1-grade estimate for the $\bar{x}$-localized covariogram, which no longer factorizes exactly; and (c)* condition (M) in full generality*. By Proposition 8.6, the $u \leftrightarrow v$ parity, the absence of quarter-power classes and the continuum constant are conditional on (M) alone — no class hypothesis beyond it — and the reach of Theorem 8.2 is governed by that same single null-displacement condition (M), which is proved for the causal diamond, for the polytopal family and for the exactly solvable smooth family, and confirmed numerically on a six-region library. What is* not *proved is (M) for every compact causally convex region: covariograms of semianalytic regions are subanalytic, and subanalyticity alone permits Puiseux exponents with denominator > 1, so a general proof must exclude a degenerate null tangency of the maximising locus (Λ tangent to the null line to higher order). No region exhibiting one is known to us, and none occurs in the library; it is not excluded a priori. This residue is strictly smaller than the question posed in earlier versions — it is a statement about one limit, $v \to 0^+$, rather than about a global power series — and two side conditions travel with it: all exact work here assumes $\hat{\omega}$-independent g_R (for anisotropic*

regions Proposition 8.6(ii),(iii) are unaffected, but the explicit reduction acquires extra terms), and the 2+1*-dimensional analogue is untreated, with logarithmic rather than half-integer interior loci expected. (An earlier version of this problem conjectured the mechanism through smoothness of* g_R *in* $(w^0, |\vec{w}\,|^2)$ *together with a rotation-invariance hypothesis; both are superseded — the operative mechanism is the* $\mu \mapsto 1/\mu$ *symmetry of the boost window, and the rotation hypothesis was never needed.) No structural obstacle is visible in (a)–(c), but (a) and (b) are real computations: the smooth-field* $3{+}1D$ *limit is not claimed.*

For the capped *observable of Section 9.4, (b) and (c) do not arise. Since* $\max(u, v) \le \ell_m$ *bounds the translate distance, the localized covariogram of any compact causally convex region is exactly constant on the bulk, so* uvH_R/π *reduces there to the degree-2 jet* $\mathrm{Vol}(R)[M_{00}(u+v)^2 + \bar{M}(u-v)^2]$ *— a polynomial, hence the easiest member of the integer-power class, with* $T_R = \ell_m$ *— and the region enters only through its volume plus a boundary layer of relative width* $O(\ell_m)$*. The delicate null-boundary and tip expansions never occur, because under the cap no loop reaches the boundary from the bulk.*

Two remarks close the section: one on where the magnetic field actually lives in this construction, and one on why no simulation will ever see the 3+1-dimensional coefficient.

Remark 8.9 (the magnetic sector rides entirely on the boost tail). By Lemma B.1, $\Sigma^{ij} \equiv 0$ per loop: *a chain-pair loop never sees the magnetic field in the rest frame of its interval.* Yet the continuum limit is exactly η-contracted — the full $F_{\mu\nu}F^{\mu\nu}$, magnetic sector included. The magnetic half of Maxwell theory is therefore generated *entirely* by the boost tail of the axis ensemble — precisely the sector controlled by the most delicate cancellation in the paper, the moment kills. (The end-to-end Monte Carlo of Appendix B, Step 4, Tier 3 tests exactly this sector — a pure-magnetic field, for which the rest-frame answer is zero — and confirms the exact mean at that test's resolving power: the standard error is 3.7% of the mean, so the boost-tail-generated magnetic sector is verified at the 4% level.) This raises a physical question that the constant-field average cannot answer: does a chain-pair action have enough support to carry a *dynamical* magnetic sector, or is its emergence an artefact of the constant-field average? A natural next object is a loop class with non-degenerate spatial projection — three intermediates (e.g. the loop $x \to z_1 \to z_2 \to y \to z_3 \to x$), or a z, z' pair not both drawn from the same interval — which would couple to $\vec{B}$ locally rather than through the boost tail.

Remark 8.10 (numerical reachability). The approach (13) is glacial, and it is honest to quantify this rather than let the word "verified" imply otherwise. At $\rho = 400$ the mean has the *wrong sign* for a pure electric field ($r_4 = +5.2 \times 10^{-3}$ against the limit $-3/(10\pi) = -0.0955$); r_4 crosses zero only near $\rho \approx 1300$; at $\rho = 10^6$ it reaches -0.0781, still 18% short. Solving (26): reaching 10% accuracy requires $\rho \approx 5 \times 10^6$ ($N \approx 6.6 \times 10^5$ elements on the unit diamond); 5% requires $\rho \approx 2.9 \times 10^7$ ($N \approx 3.9 \times 10^6$); 1% requires $\rho \approx 1.4 \times 10^9$ ($N \approx 1.8 \times 10^8$). Since sprinkling simulations are $O(N^2)$ in relations and the action is a four-point sum, no direct causal-set simulation will confirm the 3+1-dimensional coefficient. We regard this not as an embarrassment but as a physical observation about the marginality of $d = 4$: the coupling is dimensionless there, the corrections form a $\ln \rho/\sqrt{\rho}$ series with large constants, and the continuum coefficient is reachable by exact computation only — which is what Theorem 8.2 provides.

9 The four-dimensional variance: a diagonal-class divergence theorem

An earlier version of this paper listed the 3+1-dimensional variance — the analogue of Theorem 5.3, without which Theorem 8.2 is a statement about the ensemble mean only — among its open problems. The question is now answered at the level of the diagonal class, and the answer is negative: the diagonal class of the variance *diverges*, with exact law, and its structural cause shows that no redesign of the coefficients within this action class can cure it. What is proved is one-sided: the $o(\rho^{1/2})$ bound on the signed off-diagonal sharing classes is an open lemma (Section 9.2), with numerical evidence that the diagonal class dominates the total. We state the result at full strength precisely because it is unwelcome.

9.1 The diagonal class, exact

Expand $\mathbb{E}[S^2]$ by the multivariate Mecke equation into classes according to how the two interval pairs (x_1, y_1), (x_2, y_2) and their intermediates share points. The class of disjoint pairs with disjoint intervals factorizes and cancels against $\mathbb{E}[S]^2$ identically (Poisson independence); the classes sharing one or more points carry signed products $C_{n_1}C_{n_2}$; and the *diagonal* class — same interval pair $(x_1, y_1) = (x_2, y_2)$, all pairs of chain pairs within it — carries the *squares* C_n^2 and is strictly positive. It is also exactly computable, by the same covariogram identity (Lemma B.2) that collapsed the mean.

Theorem 9.1 (diagonal-class divergence, exact)**.** *Let $V_{\mathrm{diag}}(\rho)$ denote the diagonal class of* $\mathrm{Var}[S]$ *for the killed ray* (12) *on the diamond of height T, pure-electric normalization ($E = 1$). With $s = \pi\rho u^2v^2/24$, $h = \cosh^2\chi = (u+v)^2/(4uv)$, $P_{\mathrm{ang}}(h) = \frac{1}{5}(1-h)^2 + \frac{2}{3}h(1-h) + h^2$ (the S^2 average of $|E'|^4/E^4$), $\mathfrak{m}_n = \frac{1}{2}n(n-1)$, $\mathfrak{t}_n = n(n-1)(n-2)$, $\mathfrak{q}_n = \mathfrak{m}_n^2 - \mathfrak{m}_n - \mathfrak{t}_n$, and the rest-frame quartic moments $a_4 = \frac{101}{134400}$, $a_3 = \frac{101}{268800}$, $a_2 = \frac{1}{3600}$,*

$$V_{\mathrm{diag}}(\rho) = \rho^2\,\frac{\pi}{24}\,\frac{\pi}{2}\iint_{0<v<u<T} \mathrm{d}u\,\mathrm{d}v\;(u-v)^2(T-u)^2(T-v)^2\,(uv)^4\,P_{\mathrm{ang}}(h)\,K_d(s), \tag{16}$$

$$K_d(s) = e^{-s}\sum_{n=2}^{5} C_n^2\,\frac{s^n}{n!}\big[a_4\mathfrak{m}_n + a_3\mathfrak{t}_n + a_2\mathfrak{q}_n\big],$$

an exact identity at every density, and

$$V_{\mathrm{diag}}(\rho) \;\xrightarrow[\rho\to\infty]{}\; \frac{\sqrt{6\pi}}{10080}\,M_{1/2}^d\,\rho^{1/2}\,T^{10}, \qquad M_{1/2}^d = \sum_n C_n^2\big[a_4\mathfrak{m}_n + a_3\mathfrak{t}_n + a_2\mathfrak{q}_n\big]\,\frac{\Gamma(n+\frac{3}{2})}{n!} = 5.5349, \tag{17}$$

i.e. $V_{\mathrm{diag}} \to 2.3840\times10^{-3}\sqrt{\rho}$ at $T = 1$ (quadrature of (16) *at $\rho = 10^{10}$: 2.3851×10^{-3} with local log–log slope $+0.4995$). Relative to the limiting mean $\mathbb{E}[S]/(E^2V_D) \to -3/(10\pi)$, the diagonal class alone gives*

$$\frac{V_{\mathrm{diag}}(\rho)}{\mathbb{E}[S]^2} \;\longrightarrow\; 15.26\,\sqrt{\rho}\,T^2 \;=\; 42.2\,\sqrt{N}, \qquad N = \rho V_D: \tag{18}$$

the variance is super-extensive ($V_{\mathrm{diag}} \propto V_D^{5/2}$ at fixed $\rho\ell^4$), and the relative variance grows *with both density and region size.*

Proof sketch. Condition on the pair (x, y) and on $n(x, y) = n$: the intermediates are i.i.d. uniform, and the sum of Φ^2 over the $\binom{n}{2}$ chain pairs has a second moment whose z-terms reduce to three

rest-frame conditional moments of the double cone (computed from $\sigma^2 = \mathbb{E}[X^2] = \frac{1}{30}$, $\mathbb{E}[X^4] = \frac{3}{1120}$ for the axis component X at $\tau = 1$: the four-point moment $\frac{101}{8400}$, the shared-intermediate moment $\frac{101}{16800}$, and the disjoint-pair square $\frac{1}{225}$, weighted by the pair counts $\mathfrak{m}_n, \mathfrak{t}_n, \mathfrak{q}_n$), while the angular factor is the S^2 average of $|E'|^4$ with $|E'|^2 = E^2(\cos^2\alpha + \cosh^2\chi \sin^2\alpha)$. The region integral collapses by Lemma B.2 exactly as in Theorem 8.2, giving (16). For the asymptote, the kernel localizes $v \sim \rho^{-1/2}/u$ while the u-integral spreads over $u \sim O(T)$; the substitution $s = \pi\rho u^2 v^2/24$ turns (16) into $\rho^{-3/2}$ times half-integer moments of K_d, of which the leading one is $M^d_{1/2}$, with the u-integral contributing $\int_0^T u^5(T-u)^2\,\mathrm{d}u$-type factors that assemble to the displayed constant. The quadrature pipeline evaluating (16) reproduces the *mean* closed form (26) to 4×10^{-15} relative at $\rho = 400$ (2×10^{-9} at 10^7) when run on the mean kernel — the measure bookkeeping is validated at theorem level (ancillary script `p18_v2_variance.py`, stages `meancheck`, `diagquad`, `momcheck`; the conditional moments are Monte Carlo checked to $< 10^{-3}$). □

Where the divergence lives: near-null *slivers* — intervals stretched a finite fraction of the diamond along a null ray ($u \sim T$) with volume $\sim 1/\rho$ ($v \sim \rho^{-1/2}/u$), boosted so hard that $\cosh^4\chi \sim (u/v)^2$ beats both the $\sinh^2\chi$ measure and the kernel. Each sprinkling holds $\sim \rho^{3/2}$ such sparse pairs, each contributing $\sim \rho^{-1/2}$ with C_n-alternating signs: the kill cancels the *mean* of this population, but its Poisson shot noise survives. This is kin to the known fluctuation pathology of the BDG-type discrete actions, whose cure is mesoscale smearing [4, 9].

Proposition 9.2 (no coefficient ray cures the diagonal class)**.** *The bracket $a_4\mathfrak{m}_n + a_3\mathfrak{t}_n + a_2\mathfrak{q}_n$ is strictly positive for every $n \geq 2$. Hence every moment of K_d is strictly positive for every nonzero real coefficient vector of every finite support: no moment-kill exists for the diagonal class, for* any *redesign of the C_n within the chain-pair action class.*

Proof. Direct evaluation: the bracket equals 7.51×10^{-4}, 4.51×10^{-3}, 1.52×10^{-2}, 3.84×10^{-2} at $n = 2, \ldots, 5$, and is a positive combination for every n: the $\mathfrak{q}_n$ term, the only one not manifestly positive, is in fact $\mathfrak{q}_n = \mathfrak{m}_n^2 - \mathfrak{m}_n - \mathfrak{t}_n = \frac{1}{4}n(n-1)(n-2)(n-3)$ exactly — the number of *ordered* pairs of disjoint chain pairs in an n-element interval, $\binom{n}{2}\binom{n-2}{2}$, hence a nonnegative integer by its combinatorial meaning, with no estimate required. Since K_d carries C_n^2 against these weights, positivity is termwise. The structural point: every object the kill system of Sections 3–8 annihilates — divergent moments, Taylor channels, the boost channel — is *linear* in the C_n; the variance diagonal is *quadratic*, and the kill cannot propagate to squares. □

Remark 9.3 (why 1+1 dimensions escapes). The escape is not the kill. In 1+1D the loop observable A is boost invariant (Lemma 3.1), so the boost channel of the second moment contributes one logarithm — the same class is $O(\rho^{-3}\ln\rho)$, which is Theorem 5.3. In 3+1D the loop bivector $\Sigma^{0i} = \frac{\tau}{2}(z'-z)^i$ (Lemma B.1) is frame dependent, $E_n^2 \sim \cosh^2\chi$, the boost tail $\int \sinh^2\chi \cosh^4\chi\,\mathrm{d}\chi$ grows as $(T/\tau)^6$, and the two-dimensional logarithms become powers. That single structural difference is the whole divergence. In general dimension the same power counting ($uv \sim \rho^{-2/d}$, $h \sim \rho^{2/d}$) indicates a diagonal-class relative variance $\sim \rho^{2/d}$ for every $d \geq 3$ — a scaling expectation, not a result at the level of Theorem 9.1; the $d = 3$ counterpart of the exact computation has not been performed. Subject to that scaling, the 2+1-dimensional pilot of Section 7 is *also* expected to be an ensemble-level statement (relative variance $\sim \rho^{2/3}$), and only $d = 2$ self-averages.

9.2 Numerical confirmation

End-to-end Poisson sprinklings (killed action per configuration, unit diamond, pure electric background; $M = 114$–2000 sprinklings per density, $\rho = 10^3$–1.6×10^4, N up to ≈ 2100; fixed published

seeds; ancillary script `p18_v2_variance.py`, stages `sprinkle` and `fit`) give ensemble variances of $r = S/(E^2 V_D)$ of 4.37, 7.82, 13.9, 16.7, 22.0 — monotone growth with fitted exponent

$$\mathrm{Var}[r] \;\propto\; \rho^{+0.676\pm0.031}$$

over the 16× range, the finite-density face of the asymptotic $+\frac{1}{2}$ (the exact diagonal-class curve's local slope runs $0.82 \to 0.46$ over the same window). The measured total-variance to diagonal-quadrature ratios are 1.04, 1.06, 1.22, 1.02, 0.98: across the tested range the diagonal class tracks the total — the signed sharing classes are numerically subdominant, and a conspiracy in which they cancel V_{diag} is disfavored by these data. The evidential weight of five points must be stated honestly: the ratio sequence is non-monotone and dips below 1, and over the same window the exact diagonal curve is itself still in transient (local slope $0.82 \to 0.46$ against the asymptotic $\frac{1}{2}$), so these ratios are evidence of dominance, not proof. Since $\mathrm{Var}[S] \geq 0$ the sharing classes are bounded below only by $-V_{\mathrm{diag}}$, and nothing analytic in this paper excludes leading-order cancellation: the $o(\rho^{1/2})$ bound on the signed classes is not proved, and that lemma is what separates the present exact-plus-numerical statement from a two-sided theorem $\mathrm{Var}[S] = \Theta(\rho^{1/2})$ (see also Open Problem 9.10(a)). Within the same sprinklings the per-configuration diagonal-class Monte Carlo matches its quadrature at pulls $|\leq 1.4|$, and the ensemble *means* match the exact $r_4(\rho)$ curve of (26) where resolvable — the mean theorem and the variance divergence are confirmed in the same data. For contrast, the raw (C_2-only) channel self-averages without difficulty: the kill removes the mean, not the noise.

9.3 What this means

The consequence, stated plainly: **the 3+1-dimensional constant-field Maxwell result (Theorem 8.2) is to be read as an ensemble-mean theorem; per configuration, fluctuations dominate for this action class at every density — the diagonal class is proven divergent, and the measured totals track it — a single sprinkling's killed action is noise of relative standard deviation $\sim \rho^{1/4}$ riding a mean of order one — and the 1+1-dimensional L^2 theorem (Corollary 5.4) is unaffected.** Remark 8.10 is thereby upgraded: the unreachability of the coefficient by simulation is not a resolution problem but a divergence. The mean theorem itself is untouched — the ensembles above reproduce its exact curve — and the covariogram identity gains a second dividend, having collapsed the diagonal class too.

Three repair routes suggest themselves, and the natural question is whether the divergence is a property of the *substrate* or of the *observable*. *(i)* A loop class whose per-loop square is boost invariant — as the area A is in 1+1D (Lemma 3.1) — would remove the $\cosh^4$ enhancement at the root; one would build it from Lorentz invariants of the loop configuration, or from bivectors of invariant norm. *(ii)* Mesoscale smearing or cutting, the standard cure for the fluctuations of the BDG scalar action [4, 9]; note that the quantization stance of Section 10.2 already required smearing for a fluctuating connection — the present result says the background-field sector does too. *(iii)* The ensemble-level reading: accept $\mathbb{E}[S]$ as the physical object, a process-level rather than configuration-level ontology. This is coherent, and it is the reading under which every 3+1-dimensional statement of this paper stands as proved; its interpretational cost must be stated equally plainly: the action of a single causal set is then *not* the Maxwell action, and the continuum limit is a property of the sprinkling process, not of any one discrete universe.

The next subsection settles all three. Route (i) is closed *negatively*, at theorem grade and in both of its forms. Route (ii) works, and both its mean-preservation and its L^2 convergence are theorems rather than hopes — but only inside a co-scaling window, whose two edges have different characters: the upper edge $a = \frac{1}{4}$ is physical (above it the boost average no longer completes under the cap and the magnetic sector fails), while the lower edge is an artefact of proof technique, at

$a = \frac{1}{8}$ unconditionally and at $a = \frac{1}{12}$ if diagonal dominance holds under the cap. That price is stated as loudly as the cure. Route (iii) is therefore no longer forced; it remains the correct reading of the *uncapped* observable, for which Theorem 9.1 stands as the obituary.

9.4 Repairing the variance: what fails and what works

No intrinsic repair exists

The first route asks for a loop observable whose fluctuations do not see the boost. In 1+1 dimensions such an observable exists and is the whole escape (Remark 9.3). In 3+1 dimensions it does not, and the obstruction is not a failure of ingenuity but a count.

Proposition 9.4 (intrinsic invariant weights are boost-blind)**.** *Let w be any function of the Lorentz invariants of the chain-pair configuration (x, y, z, z') — a reweighting, or (taking w an indicator) a selection rule; in particular w is a function of the configuration* geometry*, independent of the field and of the holonomy datum Φ, and this hypothesis is load-bearing (see the scope remark after the proof) — and let S_w be the action* (3) *with Φ^2 replaced by $w\,\Phi^2$. Then:*

(i) the configuration carries exactly six functionally independent invariants, namely τ_{xy}^2 and the five relative invariants $(z-x)^2$, $(y-z)^2$, $(z'-x)^2$, $(y-z')^2$, $(z-z')^2$; all six are functions of the geometry of the configuration in the rest frame of the interval*, and the rapidity χ of the interval axis relative to the field frame is not among them;*

(ii) consequently the diagonal class of $\mathrm{Var}[S_w]$ is again $\Theta(\rho^{1/2}T^{10})$, with the positivity obstruction of Proposition 9.2 intact.

The candidate family of configuration-geometry weights is therefore structurally empty: no intrinsic weight and no intrinsic selection built from the configuration geometry removes the divergence.

Proof. (i) The configuration has 16 coordinates and the Poincaré group, acting freely at generic configurations, has dimension 10, so the invariant ring has $16 - 10 = 6$ functionally independent generators; the six listed are independent (their differentials are independent at a generic configuration). Each is a function of the rest-frame data alone, because *every* Alexandrov interval of $\mathbb{M}_4$ is isometric, in its own rest frame, to the round diamond of proper height τ — so at fixed τ the intrinsic configuration space is the same double cone regardless of how the interval is boosted against the lab. Sliverness is not an intrinsic property: at fixed volume, a sliver and a round interval differ by a boost, which no configuration invariant sees. (The apparently promising pair $(V, \tau^4/V)$ is degenerate: $V = \pi\tau^4/24$ identically in $\mathbb{M}_4$, so $\tau^4/V = 24/\pi$ for every interval.)

(ii) Insert w into the conditional expectation that produced (16). Since w is a function of rest-frame data, the (z, z')-average of $w^2\Phi^4$ and of the cross terms factorizes at fixed τ into w-weighted rest-frame quartic moments — the constants a_4, a_3, a_2 of Theorem 9.1 replaced by w-weighted ones — times the same angular factor $P_{\mathrm{ang}}(h)$ and the same measure. The boost integral $\int \sinh^2\chi \cosh^4\chi \,\mathrm{d}\chi$ is untouched. The w-weighted moments are again C_n^2-weighted sums of squares of rest-frame quantities, hence strictly positive, so Proposition 9.2 applies verbatim: a weight is a kernel redesign inside the same class. The growth law follows. □

Scope of Proposition 9.4. The quantification is over weights built from the configuration *geometry*. A bounded per-loop functional $f(\Phi)$ with $f(\Phi) = \Phi^2 + O(\Phi^4)$ — the compact choice $2(1-\cos\Phi)$ of Section 2 being the canonical example — is *neither* a geometry weight nor a bivector-norm reweighting (Proposition 9.5): it is a weight depending on the field/holonomy datum itself, and the factorization step in (ii) — pulling the boost factor out of the (z, z') average of $w^2\Phi^4$ — breaks

immediately if w depends on Φ. Weights depending on the interval cardinality n *are* covered (they are a C_n redesign, Proposition 9.2); weights depending on Φ are not, and "structurally empty" must not be read as covering them. For the compact action the *mean* is controlled at quartic order by Lemma 9.11 below; its variance is an open question.

The second form of route (i) is more concrete and deserves its own computation: build the loop from bivectors of *invariant* norm, in the hope that $\Sigma \cdot \Sigma$ weighting inserts sinh factors where the chain pair inserts cosh. It does insert them. It does not help.

Proposition 9.5 (invariant-norm plaquettes: cosh → sinh at zero asymptotic gain)**.** *In $d = 4$ the bivector space $\Lambda^2(\mathbb{M}_4)$ carries an invariant quadratic form of signature $(3,3)$, so $|F\cdot\Sigma|^2$ is* not *controlled by $(F\cdot F)(\Sigma\cdot\Sigma)$: no Cauchy–Schwarz bound is available. Explicitly, for loops confined to an interval whose axis has rapidity χ against a pure-electric lab field, with α the angle between the boost direction $\hat\omega$ and $\vec E$ and $h = \cosh^2\chi$, the S^2 average of the fourth power of the loop flux is proportional to*

$$\textit{electric type } (\Sigma = \tfrac{\tau}{2}\, n\wedge d): \quad P_{\rm ang}(h) = \tfrac{8}{15}h^2+\tfrac{4}{15}h+\tfrac{1}{5}, \qquad \textit{magnetic type}: \quad P^{\rm mag}_{\rm ang}(h) = \tfrac{8}{15}(h-1)^2.$$

The hoped-for sinh *factors are present — $\sinh^4\chi = (h-1)^2$ — and worthless: the two leading coefficients coincide at $8/15$. Since the divergence is carried by $h = (u+v)^2/(4uv) \sim u/4v \gg 1$, where $(1-1/h)^2 \to 1$, the diagonal class of the magnetic-type family is again $\Theta(\rho^{1/2}T^{10})$, reduced only in subleading terms; and positivity (Proposition 9.2) again blocks coefficient cures. Three-intermediate loops (Remark 8.9) carry both types and obey the same law.*

Proof. $P_{\rm ang}$ is the S^2 average of $|E'|^4/E^4$ computed in Theorem 9.1, with $|E'|^2 = E^2(\cos^2\alpha + h\sin^2\alpha)$; expanding and using $\langle\cos^4\alpha\rangle = \frac{1}{5}$, $\langle\cos^2\alpha\sin^2\alpha\rangle = \frac{2}{15}$, $\langle\sin^4\alpha\rangle = \frac{8}{15}$ gives the displayed quadratic in h. For the magnetic type, a spatial plaquette of invariant positive norm couples to the boosted magnetic field, and for a pure-electric lab field $|B'|^2 = E^2\sinh^2\chi\,\sin^2\alpha$, so the same S^2 average gives $\langle\sin^4\alpha\rangle\sinh^4\chi = \frac{8}{15}(h-1)^2$. Inserting either angular factor into (16) leaves the u- and v-densities and the kernel unchanged, and the two differ by $O(h)$ against a leading $O(h^2)$; the asymptotic evaluation of Theorem 9.1 therefore returns the same power of ρ. □

One member of the excluded-from-Proposition-9.4 family — the compact action named in Section 2, design remark (iii) — can be controlled at the level of the mean; that control is another dividend of the kill system and is recorded at the end of this section (Lemma 9.11).

Two sharpenings are worth recording, because they explain both why the escape is invisible and why 1+1 dimensions escapes.

Remark 9.6 (the escape sub-family exists but is not intrinsic; and why $d = 2$ is special)*.* Bivectors lying in the boost plane $t\wedge\hat\omega$, or fully transverse to it ($e_1\wedge e_2 \perp \hat\omega$), are boost-invariant components; a loop family whose Σ were purely of that type would couple with no enhancement at all. But $\hat\omega$ — the direction of the interval axis relative to the lab — is exactly what Proposition 9.4 shows the configuration cannot know. The escape requires frame knowledge; intrinsically it is unreachable. This also states sharply why $d = 2$ escapes: $\Lambda^2(\mathbb{M}_2)$ is one-dimensional and spanned by the invariant tensor $\epsilon_{\mu\nu}$ — the unique dimension in which no bivector direction is moved by a boost — so $\Phi = E\,A$ with A invariant (Lemma 3.1) and the Cauchy–Schwarz step is trivial. For every $d \geq 3$ the boost moves bivectors and the $\rho^{2/d}$ scaling expectation of Remark 9.3 stands.

The mesoscale cap: what works, and at what price

Route (ii) survives, and the surviving form is a cut rather than a smearing. The cut must be placed covariantly, and Proposition 9.4 says where: not in a frame the loop can manufacture — there is

none — but in the frame the *region* supplies. A compact causally convex region has an intrinsic rest frame (for the diamond D between O and P, the axis $P - O$); using it breaks Lorentz invariance at the regulator scale exactly as a lattice does, and restores it in the limit.

Definition. Let $\ell_m > 0$. For a causal pair $x \prec y$ write u, v for the null components of $y - x$ resolved along the region's axis, as in Appendix B. The *capped action* S_{ℓ_m} is (3) with the pair sum restricted to $\max(u, v) \le \ell_m$.

Proposition 9.7 (mesoscale cap: the mean is preserved exactly, the variance becomes extensive)**.** *Fix $\ell_m \in (0, T]$ and let F be constant on the diamond of height T. Then:*

(i) Mean, exact and cap-independent. $\mathbb{E}[S_{\ell_m}] \to \frac{3}{20\pi} \int_D F_{\mu\nu} F^{\mu\nu} \, \mathrm{d}^4x$ — *the same limit and the same constant as Theorem 8.2, for every ℓ_m, on the same killed ray* $(14, -16, 9, -2)$.

(ii) Variance, exact asymptote. *With $V^{\ell_m}_{\mathrm{diag}}$ the capped diagonal class,*

$$\frac{V^{\ell_m}_{\mathrm{diag}}}{V_{\mathrm{diag}}} = \frac{168}{T^8} \int_0^{\ell_m} u^5 (T - u)^2 \, \mathrm{d}u = 28 \Big(\frac{\ell_m}{T}\Big)^6 \Big[1 - \frac{12}{7} \frac{\ell_m}{T} + \frac{3}{4} \Big(\frac{\ell_m}{T}\Big)^2\Big], \tag{19}$$

so $V^{\ell_m}_{\mathrm{diag}} \propto V_D$: extensivity is restored, the super-extensive $V_D^{5/2}$ law of Theorem 9.1 is gone. On the unit diamond, for $\ell_m \ll T$, in the normalization $r = S/(E^2 V_D)$ of Section 9.2 with limiting mean $r_\infty = -3/(10\pi)$,

$$\mathrm{Var}[r] \;\longrightarrow\; 3.896 \sqrt{\rho}\, \ell_m^6, \qquad \frac{\mathrm{Var}[r]}{r_\infty^2} \;\longrightarrow\; 427 \, \frac{\sqrt{\rho}\, \ell_m^6}{T^4} \;=\; 55.9 \, \frac{(\ell_m/\ell_s)^6}{N}, \tag{20}$$

$\ell_s = \rho^{-1/4}$, $N = \rho V_D$: a $1/N$ self-averaging law with a λ^6 prefactor, $\lambda = \ell_m/\ell_s$.

(iii) Self-averaging window. *The relative variance is below unity precisely for*

$$\ell_s \;\ll\; \ell_m \;\ll\; \ell_m^* = 0.364 \, \rho^{-1/12} T^{2/3}, \tag{21}$$

a window that is nonempty and widens without bound ($\lambda \le 0.364 \, \rho^{1/6} T^{2/3}$, i.e. $\lambda \lesssim 17$ at $\rho = 10^{10}$, $T = 1$). For any co-scaling $\ell_m = \rho^{-a}$ with $\frac{1}{12} < a < \frac{1}{4}$ one has simultaneously $\mathrm{Var}[r]/r_\infty^2 = 427 \, \rho^{1/2 - 6a} \to 0$ and $\lambda = \rho^{1/4 - a} \to \infty$; the geometric-mean choice $\ell_m = \sqrt{\ell_s T}$ ($a = \frac{1}{8}$) gives $\mathrm{Var}[r]/r_\infty^2 = 427 \, \ell_s/T$.

Proof. (i) The capped mean is (14) with the datum $H_D(u, v)$ replaced by $H_D(u, v) \, \mathbf{1}\{u < \ell_m\} \mathbf{1}\{v < \ell_m\}$. The cap commutes with the S^2 average, since (u, v) are fixed under it, so the capped datum is again of integer-power class in the sense of Lemma 8.4 — with the *same* coefficients p_{ab} and the support parameter $T_R = \ell_m$ in place of T. Lemma 8.4 therefore applies verbatim: part (ii) gives the same three divergence classes $\rho \, m_0$, $\rho^{1/2} m_{1/2}$, $\ln \rho \cdot m_1$ (no quarter-power class) and hence the same kill system (12); part (iii) gives the same degree-2 jet $\mathrm{Vol}(D)[M_{00}(u + v)^2 + \bar{M}(u - v)^2]$, since the jet is taken as $u, v \to 0$, deep inside the cap. The support scale enters only through the μ-window, i.e. through $\ln(T_R/\omega)$, whose constant part $\ln \ell_m$ multiplies m_1 — which the kill annihilates. Hence the limit is exactly the limit of Lemma 8.4(iii) and is independent of ℓ_m: *the kill that makes the mean finite is also what makes it cap-independent.* The boost average completes below the cap ($\chi_{\max} \sim \ln(\ell_m/\tau) \to \infty$ whenever $\ell_m/\ell_s \to \infty$), so the magnetic sector — which by Remark 8.9 rides entirely on the boost tail — survives the regulator.

Along a co-scaled cap the finite-density error is not merely bounded but computed. On the diamond, with $J_k = \int_{\ell_m}^{T} (T-u)^2 u^{-k} \mathrm{d}u$,

$$\mathbb{E}[r_{\ell_m}] - r_\infty = \big[r_4(\rho) - r_\infty\big] + 1.02335\,\rho^{-1/2}\big[T^2 J_3 + 6T J_2 + J_1\big] + O(\rho^{-1}\ell_m^{-4}), \tag{22}$$

whose small-ℓ_m form is $0.5117\,\lambda^{-2}\big(1 + 8\,\ell_m/T + O\big((\ell_m/T)^2 \ln(T/\ell_m)\big)\big)$, $\lambda = \ell_m/\ell_s$, from the exact expansion $2\ell_m^2\big[T^2 J_3 + 6T J_2 + J_1\big] = T^4 + 8T^3\ell_m + 20T^2\ell_m^2 \ln(\ell_m/T) - 8T\ell_m^3 - \ell_m^4$ (the J_3 term contributes $-4\ell_m/T$ and the $6J_2$ term $+12\ell_m/T$ to the linear coefficient; an earlier version of this formula quoted the coefficient 12, an algebra error found by an external referee and verified symbolically here). The expansion is exact and finite: the v-degree of the integrand runs over $b = 1, \ldots, 7$, the classes $b = 1, 2, 3$ carry $m_0, m_{1/2}, m_1$ and are annihilated by the ray, and the first survivor is $b = 4$ with $m_{3/2} = -1.45397$. Hence the co-scaled mean converges for *every* $a < \frac{1}{4}$, at rate $O(\rho^{2a-1/2})$. The numerical validation of (22) is a *one-parameter* statement, not a rectangle: quadrature reproduces it along the co-scaling line $\ell_m = \rho^{-1/8}$ at ratios 0.97–1.01 over $\rho \in [10^6, 10^{10}]$, and more generally wherever $\lambda \gtrsim 5$ at the $\leq 8\%$ level; off the co-scaled family the formula degrades as $\lambda \to 1$, where the $O(\rho^{-1}\ell_m^{-4})$ remainder is no longer small — at $(\rho, \ell_m) = (10^6, 0.056)$, $\lambda = 1.8$, the exact-to-predicted ratio is 0.39.

(ii) On the ordered domain $0 < v < u < T$ of (16) the cap is $u \leq \ell_m$. In the $\sqrt{\rho}$ regime the kernel localizes $v \sim \rho^{-1/2}/u$ while the u-density is $\propto u^5(T-u)^2T^2$ (proof of Theorem 9.1), so capping simply truncates the u-integral, giving (19) against $\int_0^T u^5(T-u)^2 \mathrm{d}u = T^8/168$. Multiplying the uncapped asymptote (17) by (19), dividing by V_D^2 at $T = 1$, and keeping the leading term $28\ell_m^6$ gives the constant $28\sqrt{6\pi}M_{1/2}^d/(10080\,V_D^2) = 3.896$; division by $r_\infty^2 = (3/10\pi)^2$ gives 427, and $\sqrt{\rho}\,\ell_m^6 = \rho^{-1}(\ell_m/\ell_s)^6$ with $N = \rho V_D$ gives $427 V_D = 55.9$.

(iii) Solving $427\sqrt{\rho}\,\ell_m^6/T^4 < 1$ gives $\ell_m < 427^{-1/6}\rho^{-1/12}T^{2/3}$ and $427^{-1/6} = 0.364$; the remaining statements are substitutions. □

The upper bound that this asymptote leaves open — the sharing classes of the second Mecke moment, numerical in Panel C of Table 2 — is supplied unconditionally by the cap, at the cost of a factor λ^2 in the constant.

Proposition 9.8 (capped variance: unconditional upper bound)**.** *For $\ell_s \leq \ell_m \leq T$ and $\Lambda_F = \sup|F|$,*

$$\mathrm{Var}[S_{\ell_m}] \;\leq\; c\,\Lambda_F^4\,V_D\,\big(\sqrt{\rho}\,\ell_m^6 + \rho\,\ell_m^8\big),$$

with c depending only on the coefficients (C_n). In particular $\mathrm{Var}[S_{\ell_m}] \to 0$ along every co-scaling $\ell_m = \rho^{-a}$, $a > \frac{1}{8}$.

Proof sketch. The cap makes every admissible interval small in the *region* frame: if $\max(u, v) \leq \ell_m$ then $\operatorname{diam} I(x, y) \leq 2\ell_m$ in region coordinates, because the transverse extent is $\tau/2$ and the tips and waist rim have coordinates bounded by $\max(u, v)$. Hence $|\Phi| \leq c\Lambda_F\tau^2 \cosh\chi \leq c\Lambda_F u^{3/2}v^{1/2}$ uniformly, and the add-one cost of Lemma 5.1 obeys $|D_p S_{\ell_m}| \leq c\Lambda_F^2 \sum_{\text{active}} u_\pi^3 v_\pi$. The master corner estimates of Appendix A lift to four dimensions with one change: the splitting substitution produces $\int_0^{\ell_m} u^{1+w}\mathrm{d}u$ where in 1+1D it produced $\int \mathrm{d}U/U$ — the boost logarithm becomes a capped power. Explicitly, for $W = (u^3 v)^w$, $\rho^2 \int_{x \prec p} \int_{y \succ p} \mathbf{1}\{\text{cap}\}\, W\, g(V) \leq c\,\rho^{(1-w)/2}\ell_m^{2w+2}$, and likewise one rank down. Enumerating the sharing classes as in Appendix A, Step 4, the identical-pair classes give $\rho^{-1/2}\ell_m^6$ and the remaining classes ℓ_m^8; the Poincaré inequality then gives the displayed bound. The two terms are in ratio λ^2: the second is the cost of adding the $\sim \rho^{1/2}\ell_m^2$ active slivers through a point *coherently*, which the triangle inequality forces and the true C_n-alternating cancellation avoids. □

Corollary 9.9 (per-configuration reading, restored inside the window)**.** *Let $\ell_m(\rho) = \rho^{-a}$ with $\frac{1}{8} < a < \frac{1}{4}$. Then, for constant F on the diamond, the capped killed action of a* single *sprinkling satisfies $S_{\ell_m(\rho)} \to \frac{3}{20\pi}\int_D F_{\mu\nu}F^{\mu\nu}\,\mathrm{d}^4x$ in L^2,* unconditionally*: Proposition 9.8 bounds every second-moment class, and* (22) *controls the mean along the co-scaling. On the wider window $\frac{1}{12} < a < \frac{1}{4}$ of Proposition 9.7(iii) the same conclusion holds with the sharp constant of* (20)*, conditionally on diagonal dominance under the cap — which is numerical here (Table 2, Panel C, ratios* 1.04–1.18*) exactly as in Section 9.2.*

One dependency structure deserves to be stated explicitly, because it is the paper's best defence on the variance question: Proposition 9.8 is an *upper* bound on every capped second-moment class, so Corollary 9.9's L^2 convergence on $\frac{1}{8} < a < \frac{1}{4}$ stands *regardless* of how the uncapped off-diagonal question of Section 9.2 resolves — the constructive half of the paper does not depend on the one-sidedness of Theorem 9.1.

Numerics

Table 2 collects the three checks; all are produced and asserted by the second ancillary script (`p18_v2_repair.py`, stages `quad`, `sprinkle`, `fit`, `assertall`; fixed published seeds). Panel A exhibits the cost: at *fixed* ℓ_m the capped mean converges to the Maxwell constant — to 1.3×10^{-3} at $\rho = 10^9$, and the two caps $\ell_m = 0.3, 0.2$ agree with one another to 5×10^{-4}, which is cap-independence seen numerically — while the variance still grows as $\sqrt{\rho}$, by factors of 6–11 per hundredfold in density over the range tabulated, approaching the asymptotic factor of 10. Panel B is the co-scaled window $\ell_m = \rho^{-1/8}$: $\mathrm{Var}[r] \cdot \rho^{1/4}$ approaches the exact asymptote 3.896 from above (3.941 at $\rho = 10^{10}$, within 2%), so the variance *falls* as $\rho^{-1/4}$ where the uncapped variance of Theorem 9.1 grows as $\rho^{+1/2}$ forever; the mean runs along the same trajectory to -0.0930 against the limit -0.09549. Panel C is end-to-end: capped killed actions of real sprinklings, at the same three densities as Section 9.2, matching the capped diagonal quadrature at ratios 1.18, 1.11, 1.04 — so the capped diagonal class is again the capped variance, and no signed-class conspiracy occurs under the cap either — with absolute variances 135–191 times below the uncapped values measured in the same window. Per-configuration diagonal Monte Carlo matches its quadrature at pulls ≤ 0.9 and the means at pulls ≤ 1.5.

The measured local slope in Panel C is $+0.60 \pm 0.11$, and honesty requires saying what it is and is not: the sprinkling-accessible window sits *before* the co-scaled curve's peak near $\rho \sim 10^6$ (Panel B shows $\mathrm{Var}[r] \cdot \rho^{1/4}$ still rising at $\rho = 1.6 \times 10^4$), so the exponent measured there is not the asymptotic one. The Monte Carlo's job in Panel C is to validate the exact capped curve inside real sprinklings — which it does, at ratios ≈ 1 — and the exact curve carries the $\rho^{-1/4}$ asymptote. This is the same reachability situation as the mean itself (Remark 8.10).

The cost, stated plainly

The cure is a *window*, not a fixed-scale regulator, and the distinction is the whole honest content of this subsection. At fixed physical ℓ_m, (20) reads $\mathrm{Var}[r]/r_\infty^2 = 427\,\ell_m^6/(\ell_s^2 T^4)$, which *diverges* as $(\ell_m/\ell_s)^2$ when $\ell_s \to 0$: there is no exponent a for which the relative variance behaves as $(\ell_s/\ell_m)^a$ at fixed ℓ_m, and the reading "an effective theory of the action at a fixed mesoscale" therefore *fails*. Panel A of Table 2 exhibits this directly: the mean converges beautifully and the variance keeps growing. What works is the co-scaled mesoscale $\ell_m \propto \rho^{-a}$ — unconditionally for $\frac{1}{8} < a < \frac{1}{4}$ (Proposition 9.8), with the sharp constant on the wider $\frac{1}{12} < a < \frac{1}{4}$ — for which $\ell_m \to 0$ in physical units while $\lambda = \ell_m/\ell_s \to \infty$; then the relative variance vanishes and the mean is exact in the same limit. Two consequences deserve to be recorded on both sides of the ledger. In favour: since $\ell_m \to 0$,

A. Fixed ℓ_m (the cost): mean converges, variance still grows					
ℓ_m	ρ	mean r	Var$[r]$	small-ℓ_m asymptote / exact factor	
0.3	10^5	-0.0416	0.837	0.898	0.932
0.3	10^7	-0.0864	5.568	8.981	0.620
0.3	10^9	-0.0942	50.27	89.81	0.560 ($\to 0.553$)
0.2	10^9	-0.0939	5.567	7.884	0.706 ($\to 0.687$)
B. Co-scaled $\ell_m = \rho^{-1/8}$ (the cure): variance falls as $\rho^{-1/4}$					
ℓ_m	ρ	mean r	Var$[r]$	Var$[r] \cdot \rho^{1/4}$ (asymptote 3.896)	
0.2982	1.6×10^4	-0.0010	0.1572	1.77	
0.1778	10^6	-0.0578	0.1584	5.01	
0.1000	10^8	-0.0857	0.0444	4.44	
0.0562	10^{10}	-0.0930	0.0125	3.94	
C. End-to-end capped sprinklings ($\ell_m = \rho^{-1/8}$)					
ρ	M	Var$[r]$ (bootstrap 68%)	quadrature	ratio	uncapped (§9.2)
4000	300	0.0728 $[0.0666, 0.0790]$	0.0620	1.18	13.9
8000	150	0.1208 $[0.1071, 0.1338]$	0.1083	1.11	16.7
16000	90	0.1631 $[0.1398, 0.1822]$	0.1572	1.04	22.0

Table 2: The mesoscale cap, numerically (unit diamond, $T = 1$, pure electric background, killed ray $(14, -16, 9, -2)$; limiting mean $r_\infty = -3/(10\pi) = -0.09549$). Panel A: exact capped quadrature at fixed ℓ_m; the last column is the measured ratio of Var$[r]$ to the small-ℓ_m asymptote $3.896\sqrt{\rho}\,\ell_m^6$, converging to the exact finite-ℓ_m factor of (19) shown in parentheses. Panel B: the same quadrature along the co-scaled window. Panel C: Poisson sprinklings of the capped killed action against the capped diagonal-class quadrature; the last column repeats the uncapped variances of Section 9.2 at the same densities.

every *fixed* physical probe scale eventually sits far above the regulator, so continuum Maxwell is recovered per configuration at all fixed scales — the standard price of a regulator that vanishes in the continuum limit, and structurally the same bargain lattice field theory strikes with Lorentz invariance, except that here the frame broken at the regulator is supplied by the region's own axis rather than by a background lattice. Against: the per-configuration statement is not available at any fixed discretization of the observable; it is a statement about a family of observables indexed by ρ, and a reader who wants the action of one causal set, defined once and for all, does not have it. Route (iii) of Section 9.3 remains the honest reading of the uncapped object.

Open Problem 9.10 (residues of the repair, updated)**.** *(a)* Sharp two-sided theorem. *Proposition 9.8 bounds all classes unconditionally but is looser than the exact diagonal asymptote* (20) *by the factor $\lambda^2 = \sqrt{\rho}\,\ell_m^2$, the coherence cost of the triangle inequality on the $\sim \rho^{1/2}\ell_m^2$ slivers through a point. Recovering it — and with it the window $a > \frac{1}{12}$ and the constant* 3.896 *— requires showing that the off-diagonal sharing classes are $o(\sqrt{\rho}\,\ell_m^6)$ under the cap. The* uncapped *counterpart — an $o(\rho^{1/2})$ bound on the uncapped signed sharing classes — is the missing lemma of Section 9.2: it is what would upgrade Theorem 9.1 plus Proposition 9.2 to a two-sided theorem* $\mathrm{Var}[S] = \Theta(\rho^{1/2})$*, and neither direction of it is proved here. For the matching lower bound,* $\mathrm{Var}[S] \geq \mathrm{Cov}[S, Z]^2 / \mathrm{Var}[Z]$ *against a test statistic Z correlated with the sliver population avoids the sign problem entirely and needs one cross-moment; neither is executed here. (b)* Capped smooth-field mean. *The obstruction identified in Open Problem 8.8(a) — uniformity over the boost hyperbola — does not survive the cap:* $\max(u, v) \leq \ell_m$ *forces* $\mathrm{diam}\, I(x, y) \leq 2\ell_m$ *in region coordinates, so the Taylor remainder is uniform in χ, and each retained Taylor order is an integer-power datum whose divergences*

Lemma 8.4 assigns to the killed classes. Pricing the unstructured remainder against the absolute scale of the mean integral, $\rho^2 \iint_{\mathrm{cap}} |w_{\mathrm{mean}}||k|/V_D = 0.0129\,\rho\,\ell_m^4$, gives $\mathbb{E}[S_{\ell_m}] \to \frac{3}{20\pi}\int F^2$ for C^1 fields when $a > \frac{1}{5}$, and for C^2 fields when $a > \frac{1}{6}$ (Lemma 4.2 removing the first order). What remains is to write the four-dimensional Stokes remainder and the boundary-layer estimate with tracked constants. (c) The double limit. *Settled: see* (22).

The compact holonomy action, at quartic order

One loose end of Section 2 closes here, because it uses this section's machinery. Design remark (iii) chose Φ^2 over the compact $2(1-\cos\Phi)$; the following lemma — its argument supplied by an anonymous adversarial reviewer and verified symbolically before adoption — controls the difference at the level of the mean.

Lemma 9.11 (compact action: the quartic correction's leading divergence is killed)**.** *Let $S^{\cos}$ denote the action* (3) *with Φ^2 replaced by $2(1-\cos\Phi) = \Phi^2 - \frac{1}{12}\Phi^4 + O(\Phi^6)$, F constant on the diamond of height T. The mean of the quartic term is, up to the factor $-\frac{1}{12}$, exactly the $a_4\mathfrak{m}_n$ piece of the diagonal integral* (16) *with C_n^2 replaced by C_n: conditioning on $n(x,y) = n$, the $\binom{n}{2}$ chain pairs contribute $\mathfrak{m}_n$ equal quartic means, with the same (u,v) measure, the same u-density $u^5(T-u)^2$ and the same angular factor $P_{\mathrm{ang}}(h)$. Its leading ($\rho^{1/2}$) divergence class is therefore proportional to*

$$\sum_n C_n\,\mathfrak{m}_n\,\frac{\Gamma(n+\frac{3}{2})}{n!} \;=\; m_{1/2},$$

which the ray (12) *annihilates exactly ($m_{1/2} = 0$; asserted symbolically, stage* `moments` *of the fourth ancillary script); the subleading angular terms of P_{ang} land in the m_1 class (h-linear; killed, $m_1 = 0$) and in convergent classes (h^0), so the surviving quartic correction is $O(\rho^{-1/2})$, with first survivor $m_{3/2} = -1.453966$. Hence $\mathbb{E}[S^{\cos}] - \mathbb{E}[S] \to 0$ at quartic order and the constant-field mean limit of Theorem 8.2 is unchanged there. The scope must be stated as precisely as the result: the sextic and higher orders of the compact action would require the same accounting and are* not *performed here, so no claim of full equivalence of the two functionals is made; and no compact-action* variance *statement exists — Theorem 9.1 is proved for Φ^2, while for $2(1-\cos\Phi)$ the per-loop summand is bounded by 4 and the divergence mechanism (unbounded Φ^2 on the sliver tail) is absent by construction, which makes the compact-action variance a genuinely open — and interesting — question.*

10 Discussion

10.1 Relation to the causal-set programme

The results here are the gauge-sector counterpart of the scalar/gravity continuum-limit theorems of the BDG programme [5, 7, 8, 6, 11]: same Poisson Mecke technology, same boost noncompactness as the enemy, same layered-coefficient cure (the moment-cancellation mechanism in its general form is Aslanbeigi–Saravani–Sorkin [10]; see Remark 3.4) — and the parallel runs deeper: the BDG pointwise mean likewise converges only after Poisson-averaging over the ensemble, and the non-local d'Alembertian programme needs a mesoscale $\rho^{-1/d} \ll \ell \ll$ curvature scale for single-sprinkling variance control [4, 9, 10] — precisely the ensemble-mean-then-mesoscale-cap structure of the present paper — but a different observable class (closed-loop holonomies), for which the canonical coefficients fail and the correct kill conditions had to be derived anew, dimension by dimension ((7), (10), (12)). The exact finite-density mean (6), the annihilation Lemma 4.4, and

the L^2 convergence of Corollary 5.4 have, to our knowledge, no analogue yet in the scalar sector at this level of explicitness, and the techniques may be of use there. Gauge structure on causal sets is also of independent interest as a probe of which continuum concepts (parallel transport, flux, Wilson loops) survive the passage to a purely order-theoretic spacetime [14, 16].

10.2 Quantization stance

The moment-killed actions are *indefinite* by construction: by Corollary 3.3(i), positive coefficients are incompatible with finiteness, so alternation — hence indefiniteness of (3) as a quadratic form — is forced. This is the same Achilles heel as the BDG gravitational action, and the causal-set path-integral literature's response transfers: use the Lorentzian weight e^{iS} with analytic continuation in an auxiliary parameter, as in the 2D order MCMC work and the nonmanifoldness-suppression results [27, 28], rather than a Euclidean e^{-S} that does not exist. For the gauge action the situation is in one respect better: the continuum limit is the *Lorentzian* Maxwell action, whose indefiniteness $(E^2 - B^2)$ is physically correct, so e^{iS} needs no apology. The genuinely open problem is the measure over link phases: θ on causal relations is overcomplete (additivity along chains is a flatness constraint; deviations are the curvature degrees of freedom). A sharp first computation: for the interval ensemble with $n \in \{2, 3\}$, diagonalize the quadratic form $S[\theta]$ on relation-phase space, quotient the gauge directions $\theta_{xy} \to \theta_{xy} + \lambda_y - \lambda_x$, and determine the signature and rank of the reduced form — this decides whether Haar measure with e^{iS} yields a conditionally convergent Gaussian sector at the discrete level, and whether the moment kill survives as a property of the induced propagator.

10.3 Non-abelian outlook

Chain-pair loops have length-2 paths, so path ordering is trivial: $U_{\text{loop}} = U_{xz}U_{zy}U_{z'y}^{-1}U_{xz'}^{-1}$ is well defined without ordering ambiguity — the loop class was chosen minimal exactly for this. For a smooth background, Baker–Campbell–Hausdorff gives $U_{\text{loop}} = \exp(iF^aT^a\Sigma + O(\|A\|^2\,\text{diam}^3))$, and the quadratic term reproduces the abelian computation per generator: at leading order $\mathbb{E}[S] \to c\,\text{tr}\int F_{\mu\nu}F^{\mu\nu}$ with the *same* kill conditions. The honest caveat: the commutator corrections enter at the order at which F itself becomes A-dependent, so the leading-order statement must be checked uniformly over the boost hyperbola — the machinery of Section 4, with matrix-valued remainders; the central-symmetry Lemma 4.2 survives, being measure-theoretic. A clean pass/fail numerical test exists (constant noncommuting $SU(2)$ background on the 1+1D killed action; verify $\rho\,\mathbb{E}[S]/(\text{tr}F^2\,V_D) \to C_2/5$ with no coefficient shift) and requires nothing beyond 2×2 matrix exponentials in the shipped script's framework.

10.4 Outlook

The 1+1-dimensional theorem is, we believe, the first proof that a gauge-field action on causal sets has the correct continuum limit; the 2+1-dimensional pilot shows the mechanism is not a two-dimensional accident; and Theorem 8.2 now gives discrete Lorentz-invariant electromagnetism its continuum theorem in the physical dimension for constant field strength on a causal diamond — with the continuum coefficient, the finite-density mean, and the approach law all in closed form. The variance question that would have upgraded the 3+1-dimensional mean to a per-configuration statement is now answered, negatively at the level of the diagonal class: by Theorem 9.1 and Proposition 9.2 the diagonal class of the relative variance diverges as $\rho^{1/2}$ and no coefficient ray within this action class can cure it — numerically the diagonal class dominates the total, while the off-diagonal $o(\rho^{1/2})$ bound remains an open lemma — so Theorem 8.2 stands as an ensemble-mean

theorem — and that remains the fixed-scale statement, the one available for the observable defined once and for all. The repair question is answered too, and in both directions. There is no intrinsic cure by configuration-geometry weights or selections (Propositions 9.4–9.5): the loop configuration's six Lorentz invariants are blind to the boost, and the invariant-norm alternative trades $\cosh^4$ for $\sinh^4$ at the same leading coefficient (bounded functionals of the holonomy itself lie outside that family; for the compact action the mean is controlled at quartic order by Lemma 9.11, the variance open). There is a region-covariant one (Proposition 9.7): a mesoscale cap at $\max(u, v) \le \ell_m$ in the region's rest frame keeps the Maxwell mean exactly — the parity lemma applies verbatim with support ℓ_m — and makes the variance extensive. That statement is now two-sided enough to be a theorem: because a capped interval is small in the region frame whatever its boost, the four-dimensional add-one estimates lose their boost logarithm, every second-moment class is bounded unconditionally (Proposition 9.8), the co-scaled mean error is an exact finite moment sum (22), and per-configuration L^2 convergence holds *unconditionally* along any $\ell_m \propto \rho^{-a}$ with $\frac{1}{8} < a < \frac{1}{4}$ (Corollary 9.9); the wider window $\frac{1}{12} < a < \frac{1}{4}$ keeps the sharp constant 3.896 conditionally on diagonal dominance, and the residual gap is exactly the factor λ^2 that the triangle inequality costs on the slivers through a point. What that does *not* buy is a fixed physical mesoscale: there the relative variance still diverges, so the per-configuration reading is a property of a ρ-indexed family of observables, with the regulator vanishing in the continuum limit as on a lattice. Its residues are posed as Open Problem 9.10. The remaining 3+1-dimensional gap is the smooth-field extension. Region independence is no longer a gap but a characterization (Lemma 8.5, Proposition 8.6): the parity, the absence of quarter-power classes and the constant all rest on one null-displacement regularity condition (M) holding on every family examined, which also governs the reach of Theorem 8.2, and what survives as open — that condition in complete generality, where subanalytic Puiseux exponents are not excluded a priori — is a strictly smaller question than the one posed in earlier versions (Open Problem 8.8(c)). The quantization measure and the non-abelian escalation stand as posed above.

Data and code availability

The ancillary script `p18_maxwell_causal.py` (Python 3 with `numpy`, `scipy`, `matplotlib`; fixed published seeds) regenerates every table, figure and quoted number, and asserts the frozen ones (stage `assertall`; 33 checks, all passing). The 3+1D theorem's closed forms and their sprinkling check are asserted by the self-contained stage `s4dtheorem` (limit $= -3/(10\pi)$; exact finite-density curves for killed and raw channels against fresh Poisson sprinklings); the end-to-end Mecke Monte Carlo and the (25)-versus-(26) quadrature check of Appendix B, Step 4 (Tiers 2–3) are run by the stage `mecke4d`, which asserts the Tier-2 quadrature agreement and applies a $|\text{pull}| < 4$ regression guard to Tier 3 (an acceptance window, not an assertion of the pulls). The parity and region-independence identities of Lemma 8.4 — odd powers only, the log coefficients for a generic symmetric region factor, the diamond class table, and the assembled region-independent constant — are verified symbolically by the stage `parity` (`sympy`). A second ancillary script, `p18_v2_variance.py`, reproduces Section 9: stage `momcheck` Monte Carlo checks the rest-frame conditional moments behind (16) to $< 10^{-3}$; stage `meancheck` validates its quadrature pipeline against the exact mean (26) at machine precision; stage `diagquad` evaluates $V_{\mathrm{diag}}(\rho)$ over $\rho = 10^3$–10^{10} and asserts the asymptotic slope $+\frac{1}{2}$; stages `sprinkle/fit` regenerate the sprinkling ensembles and the fitted exponent; and its own stage `assertall` re-asserts every frozen number of Section 9 — $M^d_{1/2} = 5.5349$, the closed-form constant $\sqrt{6\pi}/10080 \cdot M^d_{1/2}$, the positivity brackets of Proposition 9.2, the measured exponent $+0.676 \pm 0.031$, and the total-to-diagonal ratios — from its results file (`p18_v2_results.json`). A

third ancillary script, `p18_v2_repair.py`, reproduces Section 9.4: stage `quad` evaluates the capped mean and capped diagonal variance at fixed and at co-scaled ℓ_m (Table 2, Panels A–B); stages `sprinkle/fit` regenerate the end-to-end capped sprinkling ensembles of Panel C together with their bootstrap intervals and fitted slope; and its own stage `assertall` re-asserts every frozen number of Section 9.4 — the degenerate invariant $\tau^4/V = 24/\pi$ and the invariant count of Proposition 9.4, the coincidence of the leading angular coefficients $8/15$ of Proposition 9.5, the closed-form constants 3.896, 427, 55.9 and $\ell_m^* = 0.364\rho^{-1/12}T^{2/3}$ of Proposition 9.7, the exact finite-ℓ_m factors (19) at $\ell_m = 0.3, 0.2$, the fixed-ℓ_m variance growth that is the stated cost, the co-scaled values $\mathrm{Var}[r]\rho^{1/4} = 5.01, 4.44, 3.94$ and their approach to 3.896, and the Panel-C variances, ratios and reduction factors — from its results file (`p18_v2_repair_results.json`). A fourth ancillary script, `p18_v2_close.py`, reproduces the closures of this version — the unconditional capped variance bound, the exact co-scaled cap tail, and the region-class characterization — and is deterministic throughout (quadrature and exact arithmetic only; no sprinklings, no seeds). Stage `moments` evaluates the half-integer moments of the killed ray and exhibits $m_{3/2} = -1.453966$ as the first survivor; stage `L3tail` checks (22) against exact quadrature along the co-scaled line $\ell_m = \rho^{-1/8}$ and across the $\lambda \geq 5$ family over $\rho \in [10^6, 10^{12}]$, carrying $(\rho, \ell_m) = (10^6, 0.056)$ as a documented off-line diagnostic where the formula does not apply (ratio 0.39); stages `L1corner` and `L1bound` measure the capped corner estimates behind Proposition 9.8 — their exponents, their analytic constants, and the coherence slack $(\mathrm{P}_c/\mathrm{P}_a)/\lambda^2 \to \kappa = 0.6124$ — and verify that the assembled bound lies above the exact capped diagonal variance at every tabulated point; stage `L2scale` measures the absolute scale $0.0129\,\rho\ell_m^4$ that prices a smooth-field remainder; stage `windows` assembles the co-scaling windows; and stage `region` verifies Lemma 8.5 and Proposition 8.6 — the cylinder counterexample, the master covariogram formula symbolically on the diamond (with its integer degrees 0–4), the perimeter coefficient $|K| = \pi T^3/6$, the ρ^0 log coefficient $-4\,g_R(0,0)$ across eight data with different supports and jets, the even-versus-odd/half parity discriminator at 50 digits over seven admissible data and one condition-(M)-violating control, and the exact support-edge reduction. Its own stage `assertall` re-asserts all of these (56 checks) from `p18_v2_close_results.json`. The four rosters together are $33 + 36 + 46 + 56 = 171$ checks, all passing. The archived outputs of full runs and all four results files (`p18_results.json`, `p18_v2_results.json`, `p18_v2_repair_results.json`, `p18_v2_close_results.json`) accompany the paper.

Acknowledgment and disclosure

The author thanks an anonymous adversarial reviewer whose observations led to the region-independence lemma (Lemma 8.4); the same reviewer's third-round observations supplied the argument of Lemma 9.11, the combinatorial identity in the proof of Proposition 9.2, and the scope corrections of this version. The author acknowledges the use of Claude (Anthropic) for assistance in research, computation, and manuscript preparation. All physical interpretations, mathematical claims, and conclusions are the author's own responsibility. This paper has not undergone external peer review.

A Proof of Lemma 5.2

Throughout, c denotes a constant depending only on the coefficients (C_n), changing from line to line, and $\Lambda = \Lambda_F$. Write $\mathcal{L} = \ln(\rho V_D)$ and abbreviate $V_1 = V_{x_1 y_1}$ etc.

Step 1: deterministic bound

By Lemma 5.1, $(D_pS)^2 \le c\,\Lambda^4\big(\sum_{\text{active}} V^2\big)^2 = c\,\Lambda^4 \sum_{\pi_1,\pi_2 \text{ active}} V_1^2 V_2^2$, a double sum over ordered pairs of active pairs $\pi_i = (x_i, y_i)$.

Step 2: joint sparsity

Each active pair carries the constraint $n_\eta(I_i \setminus \{p\}) \le 3$. For two intervals, using $\min(P_1, P_2) \le \sqrt{P_1 P_2}$ for probabilities and the Poisson tail $\Pr[n(I) \le 5] \le c\,(1+\rho V)^5 e^{-\rho V}$ (the constraint weakens by at most the two Mecke points and p),

$$\Pr\big[\text{both sparse}\big] \;\le\; c\,(1+\rho V_1)^{5/2}(1+\rho V_2)^{5/2}\, e^{-\rho(V_1+V_2)/2}.$$

The essential point is that exponential decay in *each* interval volume survives the joint constraint.

Step 3: master corner estimates

For $p \in D$ (the 1+1D diamond), $m \ge 0$, $\kappa \in (0,1]$, and $g_m(V) = V^m(1+\rho V)^a e^{-\kappa\rho V}$ with any fixed $a \ge 0$:

$$\text{(E1)}\quad \rho^2 \int_{x \prec p} \mathrm{d}x \int_{y \succ p} \mathrm{d}y\; g_m(V_{xy}) \;\le\; c\,\mathcal{L}\,\rho^{-m}, \tag{23}$$

$$\text{(E2)}\quad \rho \int_{y \succ p} \mathrm{d}y\; g_m(V_{py}) \;\le\; c\,\mathcal{L}\,\rho^{-m}. \tag{24}$$

Proof. For (E1) put $a = p - x$, $b = y - p$ (lightcone components in $\mathbb{R}^2_+$) and substitute $U = a_u + b_u$, $W = a_v + b_v$: the splitting variables integrate to at most $U \cdot W$, so the integral is bounded by $\int \mathrm{d}U\,\mathrm{d}W\,UW\,g_m(UW/2)$ over $U, W \le \sqrt{2}L$; fixing $V = UW/2$, $\int \frac{\mathrm{d}U}{U}$ over its admissible range contributes $\ln(cV_D/V)$, leaving $4\rho^2 \int_0^{V_D} V^{m+1}(1+\rho V)^a e^{-\kappa\rho V} \ln(cV_D/V)\,\mathrm{d}V \le c\,\mathcal{L}\,\rho^{-m}$ after $s = \rho V$ (the $-\ln s$ part is integrable against the exponential kernel). (E2) is the same computation one rank down. *This substitution is exactly the boost channel: the logarithm in* (E1)–(E2) *is the same noncompactness that produces the divergence of Theorem 3.2.* □

Step 4: enumeration of sharing classes

Two active pairs share 0, 1 or 2 points, and p may be an endpoint of either. In every class, bound $V_1^2V_2^2$ using AM–GM ($2V_1^2V_2^2 \le V_1^4 + V_2^4$) whenever the pairs share a point, then factor the integral using Step 2:

(a) identical pairs, p interior (2 Mecke points): (E1) with $m = 4$: $\le c\,\mathcal{L}\,\rho^{-4}$.

(b) identical pairs, p an endpoint (1 Mecke point): (E2) with $m = 4$: $\le c\,\mathcal{L}\,\rho^{-4}$.

(c) disjoint pairs (4 Mecke points): the sparsity bound of Step 2 factorizes, and the integral is the square of (E1) with $m = 2$: $\le \big(c\,\mathcal{L}\,\rho^{-2}\big)^2 = c\,\mathcal{L}^2\rho^{-4}$. This is the dominant class — the source of the $\ln^2$.

(d) one shared endpoint (3 Mecke points; e.g. $x_1 = x_2$, or $y_1 = x_2$): by AM–GM it suffices to bound $\rho^3 \int V_1^4 \sqrt{P_1}\sqrt{P_2}$; integrate the unshared point of π_2 first: $\rho \int \mathrm{d}y_2 \sqrt{P_2} \le c\,\mathcal{L}$ by (E2) with $m = 0$ (based at p or at the shared point — the estimate is uniform in the base point); the remaining factor is (E1) with $m = 4$: total $\le c\,\mathcal{L} \cdot c\,\mathcal{L}\,\rho^{-4} = c\,\mathcal{L}^2\rho^{-4}$.

(e) classes with p an endpoint of one or both pairs (1–3 Mecke points): identical treatment, replacing (E1) by (E2) for the pairs anchored at p; all give $\le c\,\mathcal{L}^2\rho^{-4}$.

Summing the finitely many classes, $\mathbb{E}[(D_pS)^2] \le c\,\Lambda^4\mathcal{L}^2\rho^{-4}$, uniformly in p (the corner estimates only improve near the boundary, where cones are truncated). ■

Theorem 5.3 follows: $\mathrm{Var}[S] \le \rho\int_D \mathbb{E}[(D_pS)^2]\,\mathrm{d}p \le c\,\Lambda^4 V_D\,\mathcal{L}^2\rho^{-3}$, and division by the squared mean $\mathbb{E}[S]^2 = \Theta(\Lambda^4 V_D^2\rho^{-2})$ gives the relative bound. The same estimates show $S \in L^2$, as the Poincaré inequality requires. The numerical check of the $\mathcal{L}^2\rho^{-4}$ law is in Figure 3.

B Proof of Theorem 8.2

Throughout, D is the diamond between O and $P = (T,0,0,0)$, $w = y - x = \tau n$ with $n = (\cosh\chi, \sinh\chi\,\hat\omega)$, $\hat\omega \in S^2$, and $u = \tau e^{\chi}$, $v = \tau e^{-\chi}$ the null components of w. We write $M_{\alpha\beta} = (F\eta F)_{\alpha\beta}$, $\bar M = \frac{1}{3}\sum_i M_{ii}$, so that $Mnn = E_n^2$ and $\mathrm{tr}_s\, M - M_{00} = F_{\mu\nu}F^{\mu\nu}$.

Step 1: the bivector identity and b_4 exact

Lemma B.1. *For the loop $x \to z \to y \to z' \to x$ with x, y on the interval axis (rest frame, proper time τ), $\Sigma^{ij} = 0$ identically and $\Sigma^{0i} = \frac{\tau}{2}\,(z'-z)^i$ exactly. Consequently, for z, z' i.i.d. uniform on the interval, $\mathbb{E}[\Sigma^{0i}\Sigma^{0j}] = \frac{\tau^2}{2}\,\mathrm{Var}[z^1]\,\delta^{ij} = \frac{\tau^4}{60}\,\delta^{ij}$ and*

$$\mathbb{E}[\Phi^2] = b_4\,V\,E_n^2, \qquad b_4 = \frac{\tau^4/60}{\pi\tau^4/24} = \frac{2}{5\pi}\,.$$

In the invariant basis the correlator is $(0,*,0)$ on $\{F^2, (F\cdot n)^2, (\tilde F\cdot n)^2\}$: the magnetic structure decouples per loop.

Proof. With x at the origin and $y = (\tau,\vec 0)$, the edge sums $\frac{1}{2}(a^\mu b^\nu - a^\nu b^\mu)$ over $x\to z\to y\to z'\to x$ collapse: edges through the origin contribute nothing, and $(z\wedge y) + (y\wedge z')$ has vanishing ij components (both spatial parts pair with $\vec y = 0$) and $0i$ component $\frac{\tau}{2}(z'-z)^i$. For the uniform measure on the double cone $\{|\vec z| \le \min(t, \tau - t)\}$ (volume $\pi\tau^4/24$): $\mathbb{E}[|\vec z|^2] = \frac{24}{\pi\tau^4}\int_0^\tau \frac{4\pi}{5}\min(t,\tau-t)^5\mathrm{d}t = \tau^2/10$, so $\mathrm{Var}[z^1] = \tau^2/30$. Then $\Phi = \frac{1}{2}F_{\mu\nu}\Sigma^{\mu\nu} = F'_{0i}\Sigma^{0i}$ (F' the field in the rest frame), and isotropy of $\mathbb{E}[\Sigma^{0i}\Sigma^{0j}]$ gives $\mathbb{E}[\Phi^2] = |\vec E'|^2\,\tau^4/60 = b_4 V E_n^2$; no component of B' enters because $\Sigma^{ij} = 0$ per loop. (The same computation in 2+1D gives $b_3 = \frac{3}{160}(12/\pi)^{4/3}$.) □

Step 2: the covariogram identity

Lemma B.2. *For future-timelike w, $D \cap (D - w) = J^+(O) \cap J^-(P - w)$ — itself an Alexandrov interval — so*

$$\int_D 1_{\{x,\, x+w\in D\}}\mathrm{d}x = \frac{\pi}{24}\Big[(P-w)^2\Big]^2 = \frac{\pi}{24}\,(T-u)^2(T-v)^2 \qquad (u < T;\ \textit{else } 0),$$

using $(P-w)^2 = T^2 - T(u+v) + uv = (T-u)(T-v)$.

Proof. $x \in D$ and $x + w \in D$ iff $x \in J^+(O)\cap J^+(O-w)\cap J^-(P)\cap J^-(P-w)$; for future w, $J^+(O-w) \supseteq J^+(O)$ and $J^-(P-w) \subseteq J^-(P)$. The volume of an Alexandrov interval of squared proper time σ^2 in $\mathbb{M}_4$ is $\pi\sigma^4/24$; here $\sigma^2 = (P-w)^2$, and $P - w$ is future-directed timelike iff $T > \tau e^{\chi} = u$. □

The future-timelike hypothesis is essential, not decorative: for *spacelike* w the nesting fails, $D\cap(D-w)$ is not an Alexandrov interval, and the volume formula is wrong (by some 24% at $w = (0.10, 0.30, 0, 0)$, $T = 1$). It holds here because $w = y - x$ with $x \prec y$ is future-timelike by

construction. The identity is the Lorentzian analogue of Matheron's set covariogram $g_K(w) = \mathrm{Vol}(K \cap (K+w))$ of Euclidean integral geometry [22, 23], for which no Minkowski counterpart appears to be in print; the nearest causal-set computations are the small-diamond chain-abundance and dimension-estimator calculations of [24, 25]. The proof is a one-line nesting argument — the identity is elementary; what it buys is that the region integral of Step 3 below is exact.

Step 3: the exact double integral and its evaluation

By the Mecke reduction (as in Theorem 3.2), Lemma B.1, Lemma B.2, the polar decomposition $\mathrm{d}^4w = \tau^3 \mathrm{d}\tau \sinh^2\chi \,\mathrm{d}\chi \,\mathrm{d}\Omega_2$, the exact sphere average $\int_{S^2} Mnn \,\mathrm{d}\Omega_2 = 4\pi[M_{00}\cosh^2\chi + \bar{M}\sinh^2\chi]$, and $\mathrm{d}u\,\mathrm{d}v = 2\tau\,\mathrm{d}\tau\,\mathrm{d}\chi$ with $\tau^2\sinh^2\chi = \frac{1}{4}(u-v)^2$, $V = \frac{\pi}{24}u^2v^2$:

$$\mathbb{E}[S] = \rho^2 b_4 \frac{\pi^3}{9216}\int_0^T\!\!\int_0^T \mathrm{d}u\,\mathrm{d}v\; uv\,(u-v)^2\big[M_{00}(u+v)^2 + \bar{M}(u-v)^2\big]\, k\Big(\frac{\pi\rho u^2v^2}{24}\Big)(T-u)^2(T-v)^2. \quad (25)$$

This is the 4-dimensional analogue of (5): an exact identity at every density. Two remarks on the order of operations, added so that the logic cannot be misread. First, the covariogram integral over x is performed *first*, and its result $\frac{\pi}{24}(T-u)^2(T-v)^2$ depends on (u,v) only — not on the spatial direction $\hat{\omega}$ of the separation; consequently the S^2 average displayed above is an exact, elementary, compact-measure average, the rapidity range is compactly cut off by the region itself ($u = \tau e^{\chi} < T$), and no regulator or non-compact axis average occurs anywhere in the proof. The exactness of the sphere average follows from the covariogram's $\hat{\omega}$-independence; Remark 8.1 is a heuristic recapitulation, not an ingredient. Second, the polar decomposition is taken with $\chi > 0$ and $\hat{\omega} \in S^2$: the prefactor is $\pi^3/4608$ on the ordered region $u > v$, written in (25) as $\pi^3/9216$ on the symmetric square (sampling $\chi \in \mathbb{R}$ together with the full S^2 would double-count each w). Substituting $u = \omega\mu$, $v = \omega/\mu$ ($\omega = \sqrt{uv}$, $\mu = \sqrt{u/v}$), the μ-integral over $(\omega/T,\, T/\omega)$ is an exact Laurent-polynomial-plus-log computation, and each resulting power ω^g maps to a kernel moment through $s = \pi\rho\omega^4/24$:

$$\int_0^T \omega^g\, k\Big(\frac{\pi\rho\omega^4}{24}\Big)\{1, \ln\omega\}\,\mathrm{d}\omega = \frac{1}{4}\Big(\frac{24}{\pi\rho}\Big)^{(g+1)/4}\Big\{m_{\frac{g-3}{4}},\ \tfrac{1}{4}\big[m^{\log}_{\frac{g-3}{4}} + m_{\frac{g-3}{4}}\ln\tfrac{24}{\pi\rho}\big]\Big\} + O\big(e^{-c\rho V_D}\mathrm{poly}\big).$$

Collecting powers of ρ gives a *finite* closed sum with the class structure

ρ^1:	$m_0\,(M_{00}+\bar{M})/2400$
$\rho^{1/2}$:	$-\sqrt{6}\,m_{1/2}\,(M_{00}+6\bar{M})/(300\sqrt{\pi})$
ρ^0:	$m_1\cdot(\ln\rho\ \text{class}) + \big[(37M_{00}-75\bar{M})m_1$
	$\quad + (6M_{00}-18\bar{M})m_1^{\log}\big]/(20\pi)$
$\rho^{-1/2}(\ln\rho)$, $\rho^{-1}(\ln\rho)$, $\rho^{-3/2}$, ρ^{-2}:	$m_{3/2}, m_2, m_{5/2}, m_3$ classes,

and *no quarter-power class*: the coefficients of $\rho^{3/4}m_{1/4}$, $\rho^{1/4}m_{3/4}$ and $\rho^{-1/4}$ cancel identically in the exact μ-integral. (This cancellation and the parity mechanism of Lemma 8.4 are one fact: the μ-range $(\omega/T,\, T/\omega)$ is symmetric under $\mu \mapsto 1/\mu$, so every monomial $u^a v^b$ survives only at the odd powers ω^{2b+1} and ω^{2a+1} — no cancellation needs to be arranged, and no smoothness in $|\vec{w}\,|^2$ is involved.) Hence $m_0 = m_{1/2} = m_1 = 0$ — the system already solved by (12) — is sufficient for every constant F, and necessary outside the measure-zero set $\{4B^2 = E^2\}$: the $\rho^{1/2}$ tensor factor is $M_{00} + 6\bar{M} = 4B^2 - E^2$, which vanishes identically for fields with $|\vec{E}| = 2|\vec{B}|$, while the m_0 factor $M_{00} + \bar{M} = \frac{2}{3}(E^2 + B^2)$ is positive for every $F \neq 0$, so $m_0 = 0$ is unconditionally necessary. On the killed ray, using $m_1 = 0$ (which also deletes every region and frame constant in the ρ^0 class)

and $m_1^{\log} = -\frac{1}{2}$:

$$\frac{\mathbb{E}[S]}{V_D} \;\to\; \frac{(6M_{00} - 18\bar{M})\, m_1^{\log}}{20\pi} = \frac{3\,(3\bar{M} - M_{00})}{20\pi} \qquad \Longrightarrow \qquad \mathbb{E}[S] \to \frac{3}{20\pi}\int_D F_{\mu\nu}F^{\mu\nu}\,\mathrm{d}^4x,$$

since $3\bar{M} - M_{00} = \mathrm{tr}_s\, M - M_{00} = F_{\mu\nu}F^{\mu\nu}$: the limit is exactly η-contracted for every constant F — the theorem-level statement of which the heuristic Remark 8.1 is a recapitulation. In the factorized form conjectured earlier, $c_4 = \kappa\, b_4 |m_1^{\log}|$ with $\kappa = \frac{3}{4}$: the S^2 average of the χ-linear coefficient of $\int E_n^2 \sinh^2\chi\,\mathrm{d}\chi$ equals $\frac{\pi}{2}F_{\mu\nu}F^{\mu\nu}$ (for pure E: $c_{\mathrm{lin}} = -\frac{1}{2} + \frac{3}{8}\sin^2\theta$, $\langle c_{\mathrm{lin}}\rangle_{S^2} = -\frac{1}{4}$). For the pure-electric normalization the full killed-ray series is

$$r_4(\rho) = -\frac{3}{10\pi} + \frac{\sqrt{6}\big[105(\ln\frac{\pi\rho}{6} + \gamma_E) - 737\big]}{32\pi\sqrt{\rho}} + \frac{24\big[36(\ln\frac{\pi\rho}{24} + \gamma_E) + 33\big]}{5\pi^2\rho} + \frac{567\sqrt{6}}{8\pi^2\rho^{3/2}} + \frac{1152}{\pi^3\rho^2}, \quad (26)$$

and for the raw (C_2-only) scheme the same computation gives $r_4^{\mathrm{raw}}(\rho) = \rho/3600 + \sqrt{6}\,\sqrt{\rho}/320 + O(\ln\rho)$: the linear divergence with exact slope $1/3600$.

Step 4: verification hierarchy

An earlier draft advertised "three independent routes" (closed form, reduced-integral Monte Carlo, sprinklings). That framing was wrong: the closed form and the reduced-integral Monte Carlo share the entire reduction of this appendix — same covariogram, same b_4, same measure, same cutoff — so an error in Lemma B.1 or Lemma B.2 would have corrupted both identically. The honest architecture is a hierarchy of four tiers, ordered by what each can and cannot test.

Tier 1: the analytic closed form (26). Because exactly seven classes arise, the series *terminates*: (26) is a finite sum, exact up to the $O(e^{-c\rho V_D}\,\mathrm{poly})$ moment-truncation tails — a strength of the covariogram route, not an asymptotic expansion.

Tier 2: quadrature of (25) *against* (26). Adaptive quadrature of the exact double integral reproduces the series to better than 10^{-13} relative at $\rho = 400$, 3820, 10^5 (script stage `mecke4d`). This checks the μ-integral algebra and the moment mapping — and nothing else: it shares every structural lemma with Tier 1.

Tier 3: an end-to-end Mecke Monte Carlo — the primary numerical evidence. We evaluate $\mathbb{E}[S] = \rho^2 \int_{x\prec y} k(\rho V_{xy})\,\overline{\mathbb{E}}_{z,z'}[\Phi^2]\,\mathrm{d}x\,\mathrm{d}y$ directly: sample x uniformly in D; sample $w = y - x$ from the polar measure $\tau^3 \sinh^2\chi\,\mathrm{d}\tau\,\mathrm{d}\chi\,\mathrm{d}\Omega_2$ with $\chi > 0$ and $\hat{\omega} \in S^2$ (taking $\chi \in \mathbb{R}$ together with the full S^2 would double-count each w); sample z, z' uniformly in $I(x, y)$; and form $\Phi = \frac{1}{2}F_{\mu\nu}(z - z')^\mu(y - x)^\nu$ from the raw loop. This uses *no* lemma of this appendix — no covariogram, no bivector identity, no b_4, no sphere average — and therefore tests the whole reduction at once. At $\rho = 400$, $T = 1$, with 2.4×10^7 samples per field sector in six independent batches (fixed published seed; script stage `mecke4d`):

field	exact (25)	end-to-end MC	pull
pure electric ($\vec{E} = \hat{x}$)	$+6.805761 \times 10^{-4}$	$+6.96(41) \times 10^{-4}$	$+0.4$
pure magnetic ($\vec{B} = \hat{x}$)	$+5.712556 \times 10^{-4}$	$+5.69(21) \times 10^{-4}$	-0.1
$\vec{E}\cdot\vec{B} \neq 0$ ($\vec{E} = \hat{x}$, $\vec{B} = (0.7, 0.4, 0)$)	$+1.051892 \times 10^{-3}$	$+0.99(6) \times 10^{-3}$	-1.1

The pure-magnetic row is the sharp test: by Lemma B.1 the loop is blind to $\vec{B}$ in the interval rest frame ($\Sigma^{ij} \equiv 0$), so for $M_{00} = 0$ the entire answer is generated by the boost tail — exactly the sector where a hidden frame artefact would live (Remark 8.9). Three qualifications keep this

table honest. (i) *It contains exactly two independent numbers.* Equation (25) depends on F only through $(M_{00}, \bar{M})$, i.e. through (E^2, B^2): the pseudoscalar invariant $\vec{E}\cdot\vec{B}$ never enters — the $(\tilde{F}\cdot n)^2$ coefficient of the loop correlator vanishes (Lemma B.1), as the parity invariance of the diamond and of the sprinkling measure requires. The third row's exact target is therefore the linear combination $\mathrm{row}_1 + 0.65 \cdot \mathrm{row}_2$ *by construction* (verified to machine precision), and its Monte Carlo content is not a third field sector but a check of the no-pseudoscalar structure: the measured combination is consistent with a vanishing $\vec{E}\cdot\vec{B}$ coefficient at -1.1σ. (ii) *Resolving power.* The pure-magnetic standard error is 3.7% of its mean, so the sharp test is a confirmation at the 4% level; the pull -0.1 should be read at that resolving power, not as sub-percent agreement. Likewise, the script's $|\mathrm{pull}| < 4$ assert is a regression guard — an acceptance window of 15–25% — not an assertion of the pulls. (iii) *Where the test is run.* $\rho = 400$ lies on the wrong-sign side of the killed channel's zero crossing (Remark 8.10), so Tier 3 deliberately operates where the mean is nowhere near its limit: it is a clean test of the exact reduction (25), not a rediscovery of the limit.

Tier 4: Poisson sprinklings. The closed form was additionally checked against (a) a Monte Carlo of the reduced mean over its five-dimensional domain (region and axis sampled, boost integrated by quadrature against the exact cutoff $\tau_m = A - \sqrt{A^2 - B}$, $A = n\cdot(P-x)$, $B = (P-x)^2$; this shares the Tier-1 reduction), and (b) fresh Poisson sprinklings:

ρ	r_4 exact (26)	r_4 reduced MC	pull	sprinklings
3.82×10^3	-0.012104	$-0.012513(467)$	-0.9	$-0.049(51)$
7.64×10^3	-0.021996	$-0.021613(919)$	$+0.4$	$-0.062(136)$
1.53×10^4	-0.032399	$-0.031405(1320)$	$+0.8$	$+0.80(93)$
10^5	-0.057779	$-0.059491(3479)$	-0.5	—
10^6	-0.078051	$-0.078527(9881)$	-0.1	—

What the sprinklings do and do not test must be stated plainly. The *raw* channel matches the exact raw curve at the 0.2–0.4% level (0.8106(18) vs 0.8091 at $\rho = 3.82 \times 10^3$; 3.983(30) vs 3.988 at 1.53×10^4) and confirms the exact slope ($r_4^{\mathrm{raw}}/\rho \to 1/3600 = 2.78 \times 10^{-4}$; measured 2.74×10^{-4} at $\rho = 3.06 \times 10^4$, still rising): this validates the measure, the covariogram and the cutoff against real sprinklings, but *not* the kill conditions — nothing is killed in the raw channel. The *killed*-channel sprinkling data (5,313 sprinklings, N up to 2000) have error bars up to 30× the signal (third row of the table): they are consistent with the prediction but are *not a test* of the killed coefficient at achievable N — and by Remark 8.10, no sprinkling simulation will be. Independently, the per-configuration limit (the boost integral at fixed $(x, \hat{\omega})$ against the exact cutoff, evaluated in 40-digit arithmetic) converges to $c_{\mathrm{lin}}(\theta)/8$ with deviation *exactly* $\propto \rho^{-1/2}$ over $\rho = 10^6$–10^{18} ($\mathrm{dev}\cdot\rho^{1/2} = 0.508, 0.496, 0.496, 0.496$) — the moment identities $\int K_1 \xi^{-2} \mathrm{d}\xi = m_0$, $\int K_1 \xi^{-3/2} \mathrm{d}\xi = 2m_{1/2}$, $\int K_1 \xi^{-1} \mathrm{d}\xi = -m_1^{\log}$ for $K_1(S) = \int_0^S sk(s)\mathrm{d}s = e^{-S}[\frac{7}{4}S^4 - \frac{5}{4}S^5 + \frac{1}{6}S^6]$ acting inside each configuration. The stage `s4dtheorem` of the ancillary script asserts the limit, the closed forms, and the sprinkling agreement; the stage `mecke4d` runs Tiers 2–3, asserting the quadrature agreement and applying the $|\mathrm{pull}| < 4$ regression guard.

References

[1] L. Bombelli, J. Lee, D. Meyer and R. D. Sorkin, "Space-time as a causal set," Phys. Rev. Lett. **59** (1987) 521.

[2] S. Surya, "The causal set approach to quantum gravity," Living Rev. Rel. **22** (2019) 5 [arXiv:1903.11544].

[3] L. Bombelli, J. Henson and R. D. Sorkin, "Discreteness without symmetry breaking: a theorem," Mod. Phys. Lett. A **24** (2009) 2579 [arXiv:gr-qc/0605006].

[4] R. D. Sorkin, "Does locality fail at intermediate length-scales?," in *Approaches to Quantum Gravity*, ed. D. Oriti (Cambridge University Press, 2009) 26 [arXiv:gr-qc/0703099].

[5] D. M. T. Benincasa and F. Dowker, "The scalar curvature of a causal set," Phys. Rev. Lett. **104** (2010) 181301 [arXiv:1001.2725].

[6] D. M. T. Benincasa, F. Dowker and B. Schmitzer, "The random discrete action for 2-dimensional spacetime," Class. Quantum Grav. **28** (2011) 105018 [arXiv:1011.5191].

[7] F. Dowker and L. Glaser, "Causal set d'Alembertians for various dimensions," Class. Quantum Grav. **30** (2013) 195016 [arXiv:1305.2588].

[8] L. Glaser, "A closed form expression for the causal set d'Alembertian," Class. Quantum Grav. **31** (2014) 095007 [arXiv:1311.1701].

[9] A. Belenchia, D. M. T. Benincasa and S. Liberati, "Nonlocal scalar quantum field theory from causal sets," JHEP **03** (2015) 036 [arXiv:1411.6513].

[10] S. Aslanbeigi, M. Saravani and R. D. Sorkin, "Generalized causal set d'Alembertians," JHEP **06** (2014) 024 [arXiv:1403.1622].

[11] L. Machet and J. Wang, "On the continuum limit of Benincasa–Dowker–Glaser causal set action," Class. Quantum Grav. **38** (2021) 015010 [arXiv:2007.13192].

[12] S. Johnston, "Particle propagators on discrete spacetime," Class. Quantum Grav. **25** (2008) 202001 [arXiv:0806.3083].

[13] S. Johnston, "Feynman propagator for a free scalar field on a causal set," Phys. Rev. Lett. **103** (2009) 180401 [arXiv:0909.0944].

[14] R. Sverdlov, "Gauge fields in causal set theory," arXiv:0807.2066 [gr-qc].

[15] R. Sverdlov, "Bosonic fields in causal set theory," arXiv:0807.4709 [physics.gen-ph]; published in Int. J. Theor. Phys. **60** (2021) 1481; see also "Quantum field theory on a causal set," PoS **LATTICE2018** (2019) 234.

[16] R. Sverdlov, "Electromagnetic Lagrangian on a causal set that resides on edges rather than points," arXiv:1805.08064 [gr-qc]; published as "Electromagnetic Lagrangian on a causal set that resides on timelike edges alone: a quasi-local model," Int. J. Theor. Phys. **61** (2022) 59.

[17] R. Sverdlov and L. Bombelli, "Gravity and matter in causal set theory," Class. Quantum Grav. **26** (2009) 075011 [arXiv:0801.0240].

[18] Nomaan X, "Quantum field theory on causal sets," in *Handbook of Quantum Gravity*, eds. C. Bambi, L. Modesto and I. Shapiro (Springer, Singapore, 2023) [arXiv:2306.04800].

[19] A. Abdi, "Quantum correlations and gauge structure from a discrete relational substrate," preprint, PhilArchive:ABDQCA (2026).

[20] F. Finster, *The Continuum Limit of Causal Fermion Systems*, Fundamental Theories of Physics **186** (Springer, 2016) [arXiv:1605.04742].

[21] N. H. Christ, R. Friedberg and T. D. Lee, "Gauge theory on a random lattice," Nucl. Phys. B **210** (1982) 310; "Random lattice field theory: general formulation," Nucl. Phys. B **202** (1982) 89; "Weights of links and plaquettes in a random lattice," Nucl. Phys. B **210** (1982) 337.

[22] G. Matheron, *Random Sets and Integral Geometry* (Wiley, New York, 1975).

[23] G. Averkov and G. Bianchi, "Confirmation of Matheron's conjecture on the covariogram of a planar convex body," J. Eur. Math. Soc. **11** (2009) 1187 [arXiv:0711.0572].

[24] M. Roy, D. Sinha and S. Surya, "The discrete geometry of a small causal diamond," Phys. Rev. D **87** (2013) 044046 [arXiv:1212.0631].

[25] D. D. Reid, "Manifold dimension of a causal set: tests in conformally flat spacetimes," Phys. Rev. D **67** (2003) 024034 [arXiv:gr-qc/0207103].

[26] G. Last and M. Penrose, *Lectures on the Poisson Process* (Cambridge University Press, 2018).

[27] S. Surya, "Evidence for a phase transition in 2D causal set quantum gravity," Class. Quantum Grav. **29** (2012) 132001 [arXiv:1110.6244].

[28] S. Loomis and S. Carlip, "Suppression of non-manifold-like sets in the causal set path integral," Class. Quantum Grav. **35** (2018) 024002 [arXiv:1709.00064].