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: on every compact causally convex region whose covariogram data admit an integer-power expansion, the same coefficient ray yields the same Maxwell limit, the region entering only through its volume. Whether every compact causally convex region admits such an expansion — null boundary segments and tips are the delicate loci — is the reduced form of the region-independence question, posed together with the smooth-field extension as open problems. 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 and variance analyses in 3+1 dimensions remain open.

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, 10]; free scalar quantum field theory on a causal set is available through Johnston's propagator constructions [11, 12].

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 [17]. The existing prior art is the corpus of Sverdlov and Bombelli: quasi-local Lagrangians for gravity and matter [16], holonomy-based Lagrangian densities for Yang–Mills fields [13, 14], and an edge-based electromagnetic Lagrangian [15]. 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 [18] 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 [20], 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 [19]; 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 across an explicitly characterized integer-power class of regions — the region enters only through its volume — reducing region independence to a single expansion question at null boundary segments and tips; that question and the smooth-field 3+1-dimensional extension remain open and are posed precisely (Open Problem 8.5).

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 [26, 27] — and the non-abelian extension are discussed, with sharply posed first computations, in Section 9.

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 $2(1 - \cos\Phi)$ (the two are numerically indistinguishable at our couplings, and Φ^2 is the one we control analytically); (iv) phases are assigned by the background-field prescription (1); dynamical link phases are beyond our scope (Section 9). 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, $V = u_0v_0/2$) 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 $t_px_q - t_qx_p = \frac{1}{2}(u_pv_q - u_qv_p)$, 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 $\mathbb{E}[A^2] = \frac{1}{16}\big[u_0^2\,\frac{v_0^2}{6} + v_0^2\,\frac{u_0^2}{6}\big] = u_0^2v_0^2/48 = V^2/12$. 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,*

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

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

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

Proof. By the multivariate Mecke equation [25], summing over the pair (x,y) and conditioning on $n(x,y) = n$ (given which the interval points are i.i.d. uniform),

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

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, $m_3 = -12 C_2$, $m_3^{\log} = m_3(\psi(6)+1)$, 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. 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 $\frac{1}{2}\big[\Phi^2 + (\Phi\circ\sigma)^2\big] = F(\bar{x})^2 A^2 + \mathcal{T}^2 + O(\|\Delta\|^3)$: 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 $|r| \leq 2V\|\partial_v F\|_\infty\sqrt{2q}$, 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

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

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 $\int_0^V |\gamma| \frac{\mathrm{d}p}{p} \leq 2c\sqrt{V}$, 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 $\mathbb{E}[\Phi^2] = V^2 \delta_{\|}(p; \bar{x}) + \varepsilon$, with $|\varepsilon| \leq 2\sqrt{\mathbb{E}[(\int h_Q F)^2] \, \mathbb{E}[r^2]} + \mathbb{E}[r^2] \leq C_F V^2 \sqrt{q}$ 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 $q = V^2/p \leq V$ on the chart, $\int_V^{p_{\max}} \sqrt{V^2/p} \, \frac{\mathrm{d}p}{p} \leq 2\sqrt{V}$, giving again the absolute $V^{5/2}$-moment, relative order $\rho^{-1/2}$; (ii) the difference channel $\gamma = \delta_{\|}(p) - \delta_{\|}(0^+)$ is $O(\sqrt{p})$ by Lemma 4.3 and is annihilated by Lemma 4.4; (iii) the constant channel $\delta_{\|}(0^+; \bar{x}) = F(\bar{x})^2/12$ 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 [25, Thm. 18.7]: $\mathrm{Var}[S] \leq \rho \int_D \mathbb{E}[(D_p S)^2] \, \mathrm{d}p$, 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 $n_\eta(I(x, y) \setminus \{p\}) \leq 3$. Then $|D_p S| \leq c \, \Lambda_F^2 \sum_{\text{active}} V_{xy}^2$, 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 $n_\eta \geq 4$ both old and new coefficients vanish (support $\{2, 3\}$), so only $n_\eta \leq 3$ contributes, with change bounded by $2 \max |C_n| \cdot \binom{4}{2} \cdot 4\Lambda_F^2 V^2$ (using $|\Phi| \leq 2\Lambda_F V$); (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 $\rho V_D \geq 2$, uniformly in $p \in D$,*

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

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 $\Lambda_F = \sup|F|$, $\rho V_D \geq 2$:*

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

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: 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 $\int F^2$: $r = 0.1893\pm0.0246$ 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 $\mathbb{E}[(D_pS)^2]$ at the diamond center falls by 4.4 orders of magnitude between $\rho = 250$ and $\rho = 4000$ while the normalized ratio $q = \rho^4\,\mathbb{E}[(D_pS)^2]/(E^4\ln^2(\rho V_D))$ 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

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).

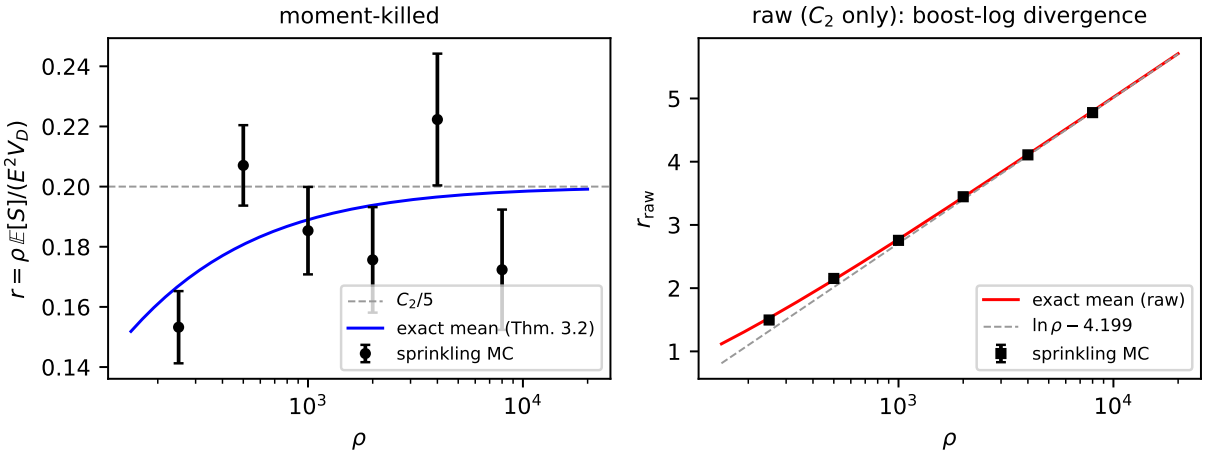


Figure 1: 1+1D convergence. Left: moment-killed normalized mean r versus density; points are sprinkling Monte Carlo, the curve is the exact finite-density mean (8) — every visible finite-size deficit is the computed boundary term. Right: the raw scheme against its exact divergent mean and the asymptote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

suggest), so quoted standard errors may be mildly optimistic; the multi-density agreement with the exact curve (8) is the load-bearing check.

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*

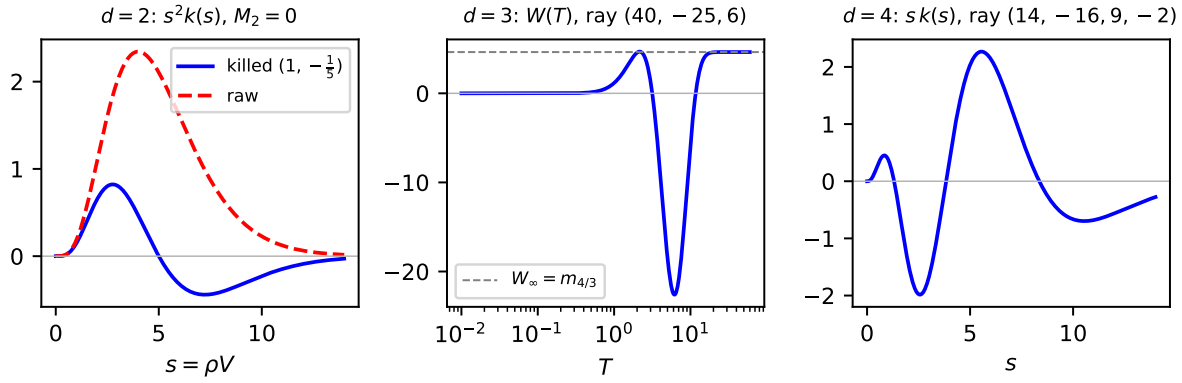


Figure 2: The kill structure in $d = 2, 3, 4$: the weighted kernels whose indicated moments must vanish. Alternation is forced (Corollary 3.3(i) and its analogues): each killed kernel must dip negative.

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 $\chi \le 1.5$), 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 8.5(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 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

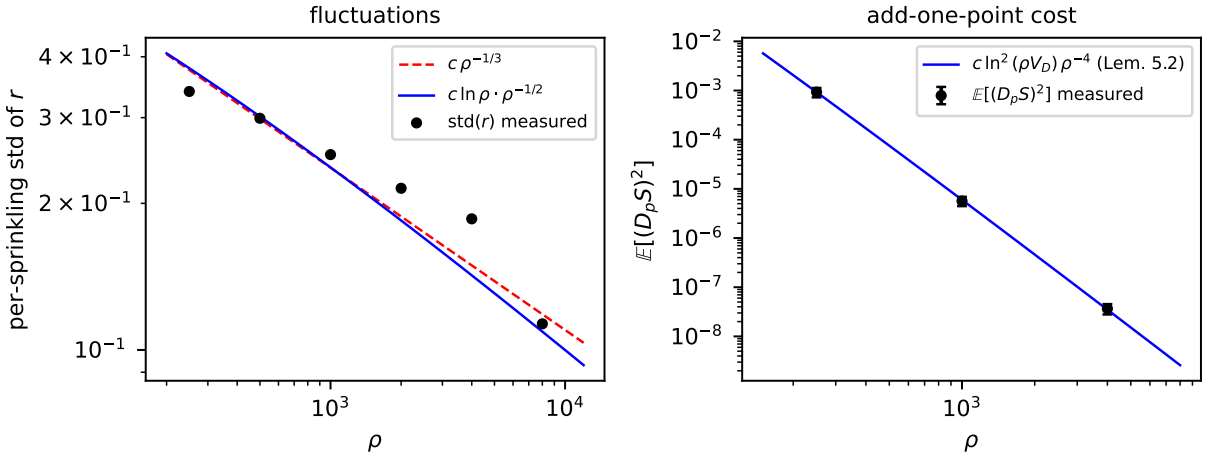


Figure 3: Left: per-sprinkling fluctuation of r; the $\ln\rho\cdot\rho^{-1/2}$ law of Theorem 5.3 and the naive $\rho^{-1/3}$ power fit are indistinguishable over this range (Remark 5.5). Right: second moment of the add-one-point cost at the diamond center versus the $\ln^2(\rho V_D)\rho^{-4}$ law of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

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 $b_4 = \mathbb{E}[(\Sigma^{0i})^2]/V = 2/(5\pi) = 0.127324$ 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).

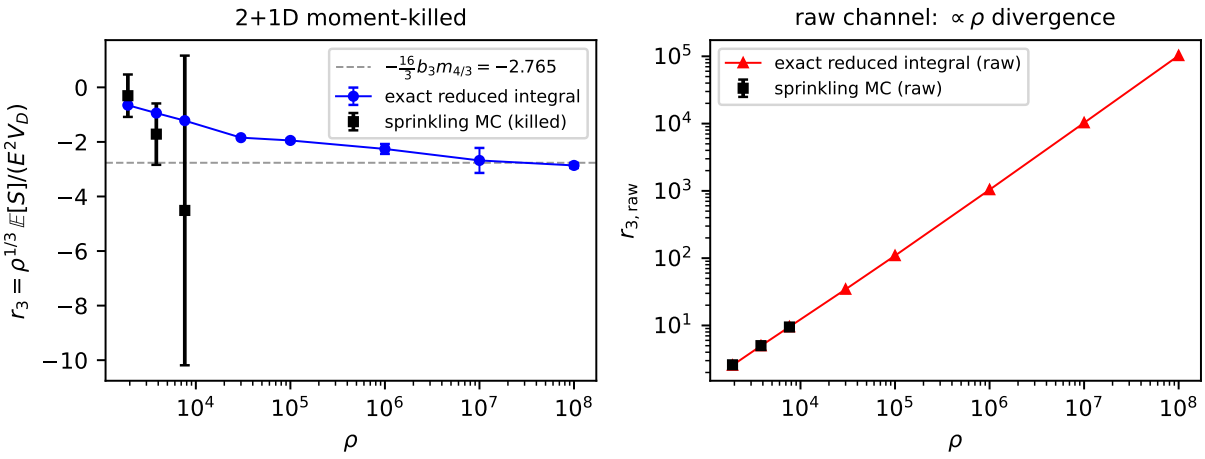


Figure 4: 2+1D pilot. Left: killed channel — exact reduced integral (circles) approaching the predicted coefficient -2.765 (dashed), with sprinkling Monte Carlo (squares). Right: raw channel — the $\propto \rho$ divergence, exact integral versus sprinkling Monte Carlo.

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.

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

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

(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.

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 [21, 22], 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* (17)–(18)*; 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*

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

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 (17)–(18) is diamond-specific: the covariogram polynomial $\frac{\pi}{24}(T-u)^2(T-v)^2$ belongs to that one region. The kill conditions and the continuum constant, however, do not: Lemma 8.4 below shows that the divergence-class structure, the kill system $m_0 = m_{1/2} = m_1 = 0$ — and with it the coefficient ray $(14, -16, 9, -2)$, which is part of the *definition* of the action — and the limiting coefficient $\frac{3}{20\pi}$ are the *same* for every compact causally convex region whose covariogram data admit an integer-power expansion (the *integer-power class* of the lemma), the region entering only through its volume. What remains region-bound is membership in that class: the diamond belongs, its covariogram being a polynomial — the easy case — and whether every compact causally convex region belongs is the single analytic question to which Open Problem 8.5(c) is now reduced; the delicate loci are null boundary segments, where the translate meets the boundary non-transversally, and the tips. The analogous statement in 2+1 dimensions is untreated. Until membership is settled in general, the result should be read as: *the action defined by the ray* $(14, -16, 9, -2)$ *has a Maxwell mean, with the same constant, on every region of the integer-power class — verified so far for the causal diamond.* 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 integer-power class, 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/π); 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. □

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

Open Problem 8.5 (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)* region independence — completion. *By Lemma 8.4, the kill system, the absence of quarter-power classes, and the continuum constant are region-independent across the integer-power class, and the $u\leftrightarrow v$ parity that drives the mechanism is automatic after the S^2 average — no rotation-invariance hypothesis is needed. What remains of (c) is a* single *analytic question:* does every compact causally convex region with piecewise-smooth boundary belong to the integer-power class — i.e. does its S^2-averaged covariogram datum uvH_R/π admit an expansion in integer powers of (u,v) on its support? *For the diamond the datum is a polynomial — the easy case. The delicate loci are null boundary segments, where the translate meets the boundary non-transversally, and the tips; the expected answer is integer-power-but-not-polynomial, giving the same three divergent classes plus a convergent tail, which is enough. (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, the smoothness route is neither necessary nor the one operating, 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, and until the membership question in (c) is settled, the theorem-grade constant-F statement is verified only on the diamond.

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 8.6 (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 8.7 (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 (18): 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 Discussion

9.1 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, 10]: same Poisson Mecke technology, same boost noncompactness as the enemy, same layered-coefficient cure — 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 [13, 15].

9.2 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 [26, 27], 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.

9.3 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 $U_{\text{loop}} = \exp(iF^aT^a\Sigma + O(\|A\|^2 \operatorname{diam}^3))$, and the quadratic term reproduces the abelian computation per generator: at leading order $\mathbb{E}[S] \to c \operatorname{tr} \int F_{\mu\nu}F^{\mu\nu}$ 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 $\rho\,\mathbb{E}[S]/(\operatorname{tr}F^2\,V_D) \to C_2/5$ with no coefficient shift) and requires nothing beyond 2×2 matrix exponentials in the shipped script's framework.

9.4 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 remaining 3+1-dimensional gaps are the smooth-field extension, the completion of region independence — Lemma 8.4 proves the kill system and the continuum constant region-independent across the integer-power class, leaving the single membership question of Open Problem 8.5(c) — and the lift of the variance theorem; 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 (17)-versus-(18) 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`). The archived output of a full run and the results file (`p18_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 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, $(D_p S)^2 \le c\,\Lambda^4 \big(\sum_{\text{active}} V^2\big)^2 = c\,\Lambda^4 \sum_{\pi_1,\pi_2 \text{ active}} V_1^2 V_2^2$, 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 $a \ge 0$:

$$\text{(E1)}\quad \rho^2 \int_{x \prec p} \mathrm{d}x \int_{y \succ p} \mathrm{d}y\; g_m(V_{xy}) \;\le\; c\,\mathcal{L}\,\rho^{-m}, \tag{15}$$

$$\text{(E2)}\quad \rho \int_{y \succ p} \mathrm{d}y\; g_m(V_{py}) \;\le\; c\,\mathcal{L}\,\rho^{-m}. \tag{16}$$

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 $V_1^2 V_2^2$ 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 [21, 22], 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 [23, 24]. 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 $\tau^2\sinh^2\chi = \frac{1}{4}(u - v)^2$, $V = \frac{\pi}{24}u^2v^2$:

$$\mathbb{E}[S] = \rho^2 b_4\,\frac{\pi^3}{9216}\int_0^T\!\!\int_0^T \mathrm{d}u\,\mathrm{d}v\; uv\,(u{-}v)^2\big[M_{00}(u{+}v)^2{+}\bar{M}(u{-}v)^2\big]\,k\Big(\frac{\pi\rho u^2v^2}{24}\Big)(T{-}u)^2(T{-}v)^2. \quad (17)$$

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 (17) 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

$$r_4(\rho) = -\frac{3}{10\pi} + \frac{\sqrt{6}\big[105(\ln\frac{\pi\rho}{6} + \gamma_E) - 737\big]}{32\pi\sqrt{\rho}} + \frac{24\big[36(\ln\frac{\pi\rho}{24} + \gamma_E) + 33\big]}{5\pi^2\rho} + \frac{567\sqrt{6}}{8\pi^2\rho^{3/2}} + \frac{1152}{\pi^3\rho^2}, \quad (18)$$

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 (18). Because exactly seven classes arise, the series *terminates*: (18) 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 (17) *against* (18). 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 (17)	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 8.6). Three qualifications keep this table honest. (i) *It contains exactly two independent numbers.* Equation (17) 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 8.7), so Tier 3 deliberately operates where the mean is nowhere near its limit: it is a clean test of the exact reduction (17), 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 (18)	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 ($r_4^{\rm raw}/\rho \to 1/3600 = 2.78 \times 10^{-4}$; 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 8.7, 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 $c_{\rm lin}(\theta)/8$ 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] 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].

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

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

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

[14] 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.

[15] 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.

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

[17] 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].

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

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

[20] 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.

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

[22] 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].

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

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

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

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

[27] 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].