Almost sure Poisson–Delaunay metric convergence

A unified proof of Theorem 1.1(i)

November 23, 2026

Abstract. Let $\mathcal{P}_I$ be the nested Poisson clouds of intensity $I\mu_h^\gamma$ obtained from a single marked Poisson process, where μ_h^γ is subcritical Gaussian multiplicative chaos on $\mathbb{R}^d$, $d \geq 2$. With the reference-metric theory specified in Section 2, a deterministic annular-median normalization of the unit-edge Euclidean Delaunay metric converges almost surely, locally uniformly, to the reference exponential metric D_h constructed from the same field. The convergence holds through every real intensity and simultaneously on every compact set. The proof develops the Poisson geometry, marked passage-law identification, scalar calibration and convergence in probability, followed by quantitative annular comparison and interpolation between the changing Delaunay graphs. The normalized metrics remain on their original field and marked-Poisson coupling.

Reference theory and status. The construction, regularity and confluence properties of the continuum reference metric are stated explicitly as background inputs. The Poisson–Delaunay argument and the almost sure upgrade are contained here. The research-proof status of the higher-dimensional background inputs is retained.

Contents

1 The model and the main result 4

2 Reference-metric hypotheses and background 5

3 Local assumptions and normalization conventions 7

4 Poisson geometry and quantitative mesh 8

5 Screened neighbors in arbitrary dimension 9

6 Mass-ball comparisons and endpoint repair 13

7 Coarse Gaussian profiles: an obstruction and a multiscale repair 15

8 A dimension-free replacement for planar joining 20

9 A past-envelope estimate for finite anchor networks 21

10 Cox–Palm screening and selected guards 27

11 The collision integral for screening 30

12 Finite-window stabilization, ties, and discrete local readouts 34

13 A mass–diameter estimate and the metric size of empty cells 38

14 Metric-ball volume inputs and Poisson counts 42

15 An almost sure hop exponent, uniformly over real intensities 44

16 Common local rules under changed field laws 48

17 Upper-order cone records without a preassigned exponent 49

18 Local transfer to a field pinned outside the buffer 51

19 Finite-projection localization of countable passage arrays 52

20 The lattice detour and full marked-path interface in every dimension 54

21 Transferring actual arrays between finite-range and pinned fields 62

22 Reference geodesic cores and local Gaussian perturbations 66

23 Gaussian comparison and local conditional laws 69

24 Quantum-label annulus transfer without a Markov decomposition 76

25 Annular comparison for passage records 83

26 Spatial amplification of a confined shortcut in every dimension 88

27 Corridor forcing and its output count 90

28 Restricted passage-law rigidity under explicit reference inputs 96

29 Propagation from selected records to every intensity 107

30 Uniform finite multiplier windows before regular variation 116

31 Confined records, escape control, and global metric convergence 118

32 GMC normalization and parameter calibration 121

33 Normalizer calibration and convergence in probability 123

34 Local tests and almost sure small-distance control 127

35 Quantitative annular coverage at finite intensity 129

36 From qualitative convergence to comparisons with constant near one 132

37 Interpolation for the marked Poisson process 137

38 Proof of Theorem 1.1(i) 139

References 141

1 The model and the main result

Fix an integer $d \geq 2$ and $0 < \gamma < \sqrt{2d}$. Let h be the coefficient-one whole-space logarithmic Gaussian field, pinned by a fixed smooth compactly supported probability density. In dimension two this is the whole-plane Gaussian free field with that normalization. We use the canonical Gaussian multiplicative chaos convention

$$\mu_h^\gamma(dx) = \lim_{\varepsilon \downarrow 0} \varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(x)}\, dx. \tag{1.1}$$

The local estimates below also use variance-normalized chaos; Section 3 records the deterministic density relating the two conventions.

Fix the deterministic scale of a reference metric D_h with the properties in Section 2. Its parameters satisfy

$$Q = Q_d(\xi) = \frac{d}{\gamma} + \frac{\gamma}{2}, \qquad \alpha = \frac{\xi}{\gamma} > 0, \qquad \Delta = \frac{\gamma}{\xi} = \frac{1}{\alpha}. \tag{1.2}$$

The reference is the subcritical Liouville quantum gravity metric when $d = 2$ and the strong exponential metric specified in Section 2 when $d \geq 3$. The field, metric and sampling measure use the same realization of h.

Conditional on h, take one Poisson process

$$\Pi \subset \mathbb{R}^d \times (0, \infty), \qquad \text{intensity } \mu_h^\gamma(dx)\, dt, \qquad \mathcal{P}_I = \{x : (x, t) \in \Pi,\ t \leq I\}. \tag{1.3}$$

The Euclidean Voronoi cells of $\mathcal{P}_I$ define its Euclidean Delaunay graph by facet adjacency. Every edge has cost one. Write $d_{\mathrm{Del},I}$ for the graph distance and $v_I(x)$ for a nearest nucleus to x, using one fixed measurable rule for ties. Define

$$d_I(x, y) = d_{\mathrm{Del},I}(v_I(x), v_I(y)).$$

These are ambient graph distances: paths may leave the compact set containing the queries. At finite intensity d_I is a pseudometric on the continuum, since two queries in the same selected cell have distance zero. The clouds are increasing in I; their graph distances need not be monotone.

Choose one deterministic annulus A whose closed observation buffer is disjoint from the support of the field pin. Let $C_I(A)$ be its minimum cell-crossing cost and $S_A(D_h)$ the corresponding reference crossing distance, as defined in Section 32. For a nonnegative random variable Z, write

$$\operatorname{med}_- Z = \inf\{t \geq 0 : \mathbb{P}[Z \leq t] \geq 1/2\}.$$

For all sufficiently large I, put

$$a(I) = \frac{\operatorname{med}_- S_A(D_h)}{\operatorname{med}_- C_I(A)}. \tag{1.4}$$

Section 33 proves that the numerator is positive, finite and a unique median, and that the denominator is positive and finite for all sufficiently large I. Define $a(I) = 1$ on the remaining bounded intensity interval.

Theorem 1.1(i) (Almost sure metric convergence)**.** *Fix d, γ and the reference metric with the properties stated in Section 2. For the marked coupling (1.3) and the deterministic normalization (1.4),*

$$\boxed{\mathbb{P}\left[\text{for every compact } K \subset \mathbb{R}^d, \quad \lim_{I \to \infty} \sup_{x,y \in K} |a(I) d_I(x, y) - D_h(x, y)| = 0\right] = 1.} \tag{1.5}$$

The limit is through every real intensity $I > 0$, on the original field and marked-Poisson coupling. In addition, a is measurable and regularly varying with index $-\alpha$:

$$a(I) = I^{-\alpha+o(1)}, \qquad \frac{a(tI)}{a(I)} \longrightarrow t^{-\alpha} \quad \text{for every fixed } t > 0. \tag{1.6}$$

Scope and proof status. The Poisson–Delaunay convergence argument is developed in this document, including the convergence in probability used in the almost sure step. The construction and stated properties of the reference metric remain background inputs, specified with their provenance in Section 2. In particular, the higher-dimensional research inputs retain their proposed-proof status. This manuscript is a research proof; no independent formal verification is claimed.

Proof organization. Sections 3–15 establish the Poisson geometry, screening, cell-size and counting estimates. Sections 16–31 construct and identify marked passage limits, propagate the identification to every intensity, and control unrestricted paths. Section 32 gives the GMC and crossing calibration. Section 33 proves the annular normalization, regular variation and convergence in probability in Proposition 33.3. Sections 34–36 turn that qualitative convergence into comparisons with polynomially small failure probabilities. Section 37 controls all intermediate Delaunay graphs between consecutive intensities of a polynomial grid. Section 38 applies Borel–Cantelli, repairs endpoints, and proves (1.5) through all real intensities on one event.

The passage-law strategy adapts the marked-path, local-kernel, annular comparison and first-failure mechanisms in [1, Sections 5–8 and 14–17], building on the field-perturbation and geodesic-saving method of [4, Sections 1.5 and 3–5]. The adapted arguments and the geometric and Gaussian estimates needed here are included below; the planar-map convergence theorem is not a convergence input. The near-one comparison argument in Section 36 is related to the LFPP argument of Devlin [16].

2 Reference-metric hypotheses and background

This section specifies the reference theory used by the graph proof. It includes the metric assignment, regularity, volume bounds and deterministic-anchor geometry. Their construction, uniqueness and confluence proofs are background results and are not reproduced here. All graph estimates and the convergence upgrade are proved from this interface in the subsequent sections.

The reference assignment. For $d \geq 3$, the assignment is specified by [14, Definition 1.1 and Theorem 1.2, p. 2]. Its subcritical condition applies because

$$Q - \sqrt{2d} = \frac{(\sqrt{2d} - \gamma)^2}{2\gamma} > 0.$$

For the logarithmic field plus a possibly dependent random continuous function, there is a measurable assignment $g \mapsto D_g$. It is a finite continuous length metric with the Euclidean topology. Its internal metrics on deterministic open sets are locally determined by g. Section 16 fixes common Borel versions under changed local field laws. The assignment has simultaneous continuous Weyl scaling, and similarity covariance for each fixed deterministic chart $x \mapsto z + rx$:

$$D_{g+f} = e^{\xi f} \cdot D_g, \qquad D_g(z + rx, z + ry) = D_{g(z+r\cdot)+Q\log r}(x, y). \tag{2.1}$$

Here the weighted length metric $e^{\xi f} \cdot D_g$ is obtained by integrating $e^{\xi f}$ against D_g-arclength and then minimizing over paths. For $d = 2$, these are the planar metric properties in [4, Theorem 1.2], also recorded in [11, Proposition 7].

The canonical pinned whole-space metric is proper and geodesic; the higher-dimensional input is [14, Lemma 2.7], and the planar properties are recorded in [4, Section 1.5]. Bounded continuous shifts preserve properness by the two-sided Weyl comparison. Only local properties are used for arbitrary unbounded shifts. Changing from the unit-sphere pin in those sources to the smooth pin changes the field by a random additive constant and is handled by simultaneous constant Weyl scaling. A prescribed countable collection of deterministic charts can be placed on one event by intersection. No simultaneous event for every uncountable choice of chart is required.

Regularity, shell costs and volume. On every compact buffer, D_h has a deterministic positive local Euclidean upper Hölder exponent: for some $\sigma > 0$ and a finite random constant C,

$$D_h(x, y) \le C|x - y|^\sigma$$

for x, y in that buffer. For $d \ge 3$, this follows from [14, Lemma 2.6]; the quantitative weak-metric estimates used below are [8, Theorem 1.7 and Propositions 4.7, 4.10 and 4.11]. Section 13 states the fixed-shell crossing and connection moment bounds at their point of use and proves the additional mass–diameter and Poisson-cell consequences.

For every compact K and every $u > 0$, almost surely, for all sufficiently small $r > 0$, uniformly over $x \in K$,

$$r^{\Delta+u} \le \mu_h^\gamma(B_{D_h}(x, r)) \le r^{\Delta-u}. \tag{2.2}$$

For $d \ge 3$, this is the input [2, Proposition 3.1, p. 9]. For $d = 2$, use [6, Theorem 1.1], absorbing its finite random constants by first using $u/2$. A constant pin change multiplies distances by $e^{\xi c}$ and mass by $e^{\gamma c}$, and therefore preserves these eventual powers because $\Delta = \gamma/\xi$. No tail estimate for the random onset scale is inferred from (2.2).

Deterministic anchors and local cores. For each fixed deterministic endpoint pair the geodesic is unique, by [2, Theorem 1.2, p. 3] when $d \ge 3$ and [5, Theorem 1.2] when $d = 2$. The graph proof uses the precise deterministic-anchor property stated in Lemma 22.3. In dimensions at least three, the constructive countable coverage in [3, Lemma 3.2, Proposition 5.1 and Sections 5.7, 7.2–7.5], together with the local escape check in Section 22, supplies that property. In dimension two, the simultaneous nearby-endpoint containment of [11, Proposition 6] gives the same interface. These source checks are the dimension-dependent inputs; saturation, passage rigidity and propagation then use a common argument.

The simultaneous anchor coverage is applied after extracting graph limits, because each extraction retains the actual fixed reference law. A contact theorem for a varying family of reference metrics is not assumed. Countable intersections over deterministic anchors and buffers are sufficient.

Other fixed weak reference metrics. The weak uniqueness inputs are [14, Theorem 1.4] for $d \ge 3$ and [4, Theorem 1.8] for $d = 2$. Thus any other fixed weak assignment is $\widetilde{D}_h = cD_h$ for a deterministic $c \in (0, \infty)$. Internal distances and cell diameters multiply by c, and

$$B_{\widetilde{D}_h}(x, r) = B_{D_h}(x, r/c).$$

Applying an estimate with half its exponent loss absorbs this constant at small radii or large intensities. This gives the stated weak-reference scope of the cell and hop estimates in Sections 13–15 and Lemma 17.4; no estimate uniform over different assignments is claimed.

The LFPP identification belongs to this background metric theory. Statements in the cited research index [15] about formalization are provenance information and are not a substitute for a proof of any input. The present graph argument does not independently replay those formalizations.

3 Local assumptions and normalization conventions

For the geometric and screening estimates, work on nested compact buffers in an open subset of $\mathbb{R}^d$. Assume that h is an actual Gaussian generalized field admitting the usual subcritical GMC construction, with local covariance

$$K(x,y) = \log \frac{1}{|x-y|} + R(x,y), \qquad |R(x,y)| \le C. \tag{3.1}$$

The remainder bound is uniform on one surrounding buffer. The regularizations satisfy the standard covariance and GMC convergence hypotheses. An arbitrary bounded function R need not define a Gaussian covariance; no such assertion is made.

Set $g = \gamma^2 < 2d$. In these local estimates, use variance-normalized GMC. We use its atomlessness, full support, local finiteness, multipoint Palm identity, Kahane comparison, positive moments below $2d/g$, and negative moments of every order for a fixed reference ball; see [9, Theorems 2.1, 2.6, 2.11 and 2.12] and [10, Propositions 3.5–3.7]. A deterministic sampling density w changes only the constants in these estimates provided that

$$0 < w_- \le w(x) \le w_+ < \infty$$

on the surrounding buffer. Positivity and an upper bound without a positive lower bound do not suffice. A fixed smooth bounded shift $h \mapsto h + f$ gives $w = e^{\gamma f}$ and satisfies these two-sided bounds on compact buffers. For a random smooth shift this observation is pathwise; an annealed estimate must also control the resulting random constants.

The canonical chaos convention used for similarity covariance is

$$\mu^{\gamma}_{h,\mathrm{can}}(dx) = \lim_{\varepsilon \downarrow 0} \varepsilon^{g/2} e^{\gamma h_\varepsilon(x)}\, dx.$$

If $\operatorname{Var} h_\varepsilon(x) = \log(1/\varepsilon) + v(x) + o(1)$, variance-normalized chaos equals $e^{-gv(x)/2}\mu^{\gamma}_{h,\mathrm{can}}$. Changing conventions can therefore insert a deterministic spatial density. When v is bounded on the surrounding buffer, as for the canonical covariance conventions used here, this density has the required two-sided bounds. The change is harmless for the local estimates, but must be included in the field when identifying the precise limiting metric. Section 32 makes this distinction explicit.

Use the marked coupling (1.3): conditional on μ, take a Poisson process on space times $[0,\infty)$ of intensity $\mu(dx)\,du$, and retain the sites with marks at most λ. Thus $\mathcal{P}_s \subset \mathcal{P}_t$ for $s \le t$. Hop distances are not asserted to be monotone in this coupling.

4 Poisson geometry and quantitative mesh

Lemma 4.1 (Uniform local filling)**.** *Let μ be a nonatomic full-support Radon measure. For every compact K and every fixed $r > 0$, almost surely all radius-r balls with centers in K contain points of $\mathcal{P}_\lambda$ for all sufficiently large λ. Empty balls whose closures meet K have radii tending uniformly to zero. Consequently Voronoi cells meeting K and Delaunay edges whose segments meet K have uniformly vanishing Euclidean diameter.*

Proof. A finite cover of K by smaller positive-mass balls is filled after a finite maximum of Poisson arrival times. If a ball of radius at least r has closure meeting K, moving a short distance toward its center produces a radius-cr subball contained in it with center in a fixed neighborhood of K. This contradicts filling if the original ball is empty.

For a Voronoi cell meeting K at z, filling first bounds the distance from z to its nucleus p. Any y in the same cell is the center of an empty ball with p on its boundary. The preceding empty-ball observation bounds $|y - p|$ uniformly. An edge has an empty witness ball containing its segment, which gives the final assertion. □

Lemma 4.2 (Mass of empty balls)**.** *For a fixed bounded buffer U, put $M = \mu(U)$. Conditional on μ, for $\lambda t \geq 8$,*

$$\mathbb{P}\big[\exists B :\ \mathcal{P}_\lambda \cap B \cap U = \varnothing,\ \mu(B \cap U) \geq t \mid \mu\big] \leq C_d(1 + \lambda M)^{d+1} e^{-c\lambda t},$$

where B ranges over Euclidean open balls. Hence, in the marked coupling,

$$\sup_{B \subset U:\, B \cap \mathcal{P}_\lambda = \varnothing} \mu(B) \leq C \frac{\log \lambda}{\lambda} \tag{4.1}$$

eventually almost surely, with a suitable fixed constant C.

Proof. Add an independent Poisson process of the same intensity. Given a witness ball of mass at least t, the second process has at least $\lambda t/2$ points in it with probability at least $1/2$. Conditional on the union process, colors are independent fair coins. Balls have VC dimension $d + 1$, so their traces on N points number at most $C_d(N + 1)^{d+1}$. A trace of size $\lambda t/2$ is entirely the second color with probability at most $2^{-\lambda t/2}$. Average over $N \sim \text{Poisson}(2\lambda M)$, absorbing constants. Conditional Borel–Cantelli on dyadic intensities, followed by monotonicity of emptiness between consecutive dyadic intensities, proves (4.1). □

Every bounded Voronoi cell can be covered by at most 5^d closed empty balls. To see this, take a $1/2$-net of directions on the unit sphere. In each associated cap, choose a farthest cell point y from the nucleus p. If x is in that cap, set $s = |x - p| \leq |y - p| = R$. The angle θ between $x - p$ and $y - p$ is at most $\pi/3$, so

$$|x - y|^2 = s^2 + R^2 - 2sR\cos\theta \leq R^2.$$

The ball $\overline{B}(y, R)$ covers that part of the cell and has empty interior.

Lemma 12.4 gives simultaneous sphere-null GMC mass and the required Delaunay general position. Thus the closed empty-ball cover and Lemmas 4.1–4.2 imply

$$\sup_{p:\, V_p^\lambda \cap K \neq \varnothing} \mu(V_p^\lambda) \leq C_d' \frac{\log \lambda}{\lambda}$$

eventually, for K compactly inside the buffer.

Proposition 4.3 (Uniform GMC mass powers)**.** *Define*

$$\beta_\pm = d + \frac{g}{2} \pm \sqrt{2dg} = \left(\sqrt{d} \pm \frac{\gamma}{\sqrt{2}}\right)^2.$$

For every sufficiently small $\varepsilon > 0$, on each compact buffer, almost surely for all sufficiently small r, uniformly in centers,

$$r^{\beta_+ + \varepsilon} \le \mu(B_r(x)) \le r^{\beta_- - \varepsilon}. \tag{4.2}$$

Fixed multiplicative constants can be absorbed by slightly reducing ε before choosing the final random threshold.

Proof. The structure-function estimate of [10, Proposition 3.7], transported by covariance comparison to the fixed buffer, gives

$$\mathbb{E}[\mu(B_r(x))^q] \le C_q r^{\zeta(q)}, \qquad \zeta(q) = \left(d + \frac{g}{2}\right) q - \frac{g}{2} q^2,$$

for $1 < q < 2d/g$ and for $q < 0$. On a grid of $O(r^{-d})$ points, Markov's inequality gives the upper-mass failure exponent $\zeta(q) - q\beta - d$. Choose $q = \sqrt{2d/g}$, which lies strictly between 1 and $2d/g$. This exponent is positive for $\beta < \beta_-$. For the lower-mass bound, choose $q = -p$ with $p = \sqrt{2d/g}$. The failure exponent is

$$p\beta - p(d + g/2) - gp^2/2 - d > 0 \quad \text{when } \beta > \beta_+.$$

Sum along dyadic radii, use containing and contained grid balls, and interpolate by monotonicity in r. □

Corollary 4.4 (Polynomial mesh and a coarse lower exponent)**.** *For every $\eta > 0$, the largest Voronoi-cell diameter meeting a fixed compact, and the largest Delaunay edge whose segment meets that compact, are eventually at most*

$$\lambda^{-1/\beta_+ + \eta}$$

almost surely. For fixed distinct x, y,

$$\liminf_{\lambda \to \infty} \frac{\log d_\lambda(x, y)}{\log \lambda} \ge \frac{1}{\beta_+}. \tag{4.3}$$

Proof. Take $r = \lambda^{-1/(\beta_+ + 2\varepsilon)}$. The lower mass bound makes the Poisson mean in each radius-cr covering ball at least a positive power of λ. There are only polynomially many such balls; their conditional filling failure is summable along dyadic intensities. Lemma 4.1's deterministic argument gives the cell and edge bounds. For (4.3), stop any path from x to y when it first leaves a fixed Euclidean neighborhood of x disjoint from y. It must travel a fixed positive Euclidean displacement in edges of the stated maximum length. No compact-containment assumption on the entire graph geodesic is needed. □

5 Screened neighbors in arbitrary dimension

Fix nested buffers and a small screening radius $\rho > 0$. For an inserted root p, let $L_{s,t}(p)$ count sites $q \neq p$ of $\mathcal{P}_t \cap B_\rho(p)$ for which some Euclidean open ball through p, q contains no point of the lower background cloud $\mathcal{P}_s$. The root is removed from the background when appropriate. The witness ball may extend beyond the buffer; the construction below extracts small empty subballs inside it and inside the buffer.

Theorem 5.1 (All finite screening moments)**.** *Under the assumptions of Section 3, including the two-sided local bounds $0 < w_- \le w \le w_+ < \infty$ if a sampling density is present, for every finite $q \ge 1$,*

$$\sup_{p\in K} \mathbb{E}[L_{s,t}(p)^q] \le C_q (t/s)^q, \qquad 0 < s \le t. \tag{5.1}$$

The estimate holds both for a deterministic query root with the original field law and for an annealed GMC Palm root. In particular,

$$\mathbb{E} \sum_{p\in\mathcal{P}_t\cap K} L_{s,t}(p)^q \le C_q t (t/s)^q. \tag{5.2}$$

Constants depend on q, d, γ, the buffers, the covariance-remainder bound, and the two density bounds when present.

The proof uses a finite cluster decomposition of the insertion positions, Gaussian covariance comparison, and an occupation-density estimate. These arguments apply in every $d \ge 2$ and are given in full below.

5.1 Finite witness choices and separation

Fix $k \ge 1$ and $n = k + 1$ insertion positions $z_0 = 0, z_1, \dots, z_k$ with

$$c_0 \le \max_i |z_i| \le C_0.$$

For a cluster I, let $r(I)$ be its connection threshold, or equivalently the largest edge length of a minimum spanning tree. Thus

$$r(I) \le \operatorname{diam}(I) \le (|I| - 1) r(I).$$

Choose a large constant $L = L(n)$. The children of an internal cluster of scale r_v are its connected components at threshold r_v/L. Recurse on nonsingleton children. For two leaves whose least common ancestor is v,

$$r_v/L < |z_i - z_j| \le (n-1) r_v. \tag{5.3}$$

There are at most $n - 1$ internal nodes and the root scale is bounded above and below.

At node v, choose a point q_v of that cluster maximizing distance to 0. It is nonzero and $|q_v| \ge r_v/2$. Its empty witness ball has radius at least $r_v/4$. Take $\delta = 1/100$ and a fixed 1/8-net of directions. A direction u within 1/8 of the inward normal at q_v gives

$$\overline{B}(q_v + \delta r_v u, \delta r_v/4)$$

strictly inside the witness ball, by the cosine rule. Put $b = \delta/[64(n+1)]$. In that ball, place $n + 1$ deterministic closed balls of radius $2br_v$, with centers separated by $8br_v$ along a coordinate direction. At most n contain inserted points. Retain the concentric radius-br_v half-ball of a ball containing none. Denote it by B_v. Then

$$\operatorname{dist}(B_v, \{z_0, \dots, z_k\}) \ge br_v, \qquad B_v \subset B(q_v, 2\delta r_v). \tag{5.4}$$

The possible choices are finite and measurable, depending only on the insertion positions. An upper bound may sum over all valid choices.

Choose $L \ge 4\delta/b$. An ancestor's ball is separated from all descendant balls by (5.4). For incomparable nodes, (5.3) and their displacements at most $2\delta r_v$ give

$$\operatorname{dist}(B_v, B_w) \ge c_n r_{\mathrm{lca}(v,w)}. \tag{5.5}$$

Indeed, in the incomparable case the lower bound is $r_{\mathrm{lca}}/(2L)$; in the ancestor case it is $br_v/2$.

5.2 The inverse-mass integral

Let $e = 1$ for a Palm root and $e = 0$ for a deterministic query. Set $\chi_0 = e$, $\chi_i = 1$ for $i \geq 1$, and use a fixed valid reference log field F on the normalized buffer. Define

$$W(z) = \prod_{0 \leq i < j \leq k} |z_i - z_j|^{-g\chi_i\chi_j}, \qquad M_z = \int_{\bigcup_v B_v} \prod_{i=0}^{k} |x - z_i|^{-g\chi_i} \mu_F(dx).$$

For each valid finite ball-choice family we claim

$$\int_{c_0 \leq \max |z_i| \leq C_0} W(z)\mathbb{E}[M_z^{-k}]\, dz_1 \cdots dz_k < \infty. \tag{5.6}$$

Put $c = d/g + 1/2 > 1$, and let $c_v = \sum_{i \in I_v} \chi_i$ be the charge of a cluster. At proper internal nodes set $P_v = (c_v - c)_+$, at the root set $P_{\text{root}} = k$, and at leaves set $P_v = 0$. The numbers

$$s_v = P_v - \sum_{w \text{ child of } v} P_w$$

are nonnegative and sum to k. This follows from the superadditivity of $(a - c)_+$ under addition of nonnegative charges; at the root the total charge is $k + e$ and $c > 1$.

The insertion weight on B_v is bounded below by a constant times

$$A_v = \prod_{u \neq \text{root} \,:\, u \text{ ancestor of } v \text{ or } u = v} (r_u / r_{\text{par}(u)})^{-gc_u}.$$

The disjointness of the balls and weighted AM–GM imply

$$M_z^{-k} \leq C \prod_v [A_v \mu_F(B_v)]^{-s_v}. \tag{5.7}$$

Dominate F on this disjoint union by a branching Gaussian field. Its common/tree variables satisfy

$$\text{Cov}(G_v, G_w) = \log(C / r_{\text{lca}(v,w)}),$$

obtained from a common root Gaussian and independent increments of variance $\log(r_{\text{par}(v)}/r_v)$. Add independent rescaled reference fields on the individual balls. By (5.5), a sufficiently large fixed C dominates the original covariance both within and across balls. This is a genuine positive semidefinite Gaussian construction, not an assumption that a pure logarithm is a covariance on an arbitrary domain.

Product negative-moment comparison follows from

$$x^{-s} = \frac{1}{\Gamma(s)} \int_0^\infty t^{s-1} e^{-tx}\, dt$$

and scalar Kahane comparison for the convex void functional with base weight $\sum_v t_v \mathbf{1}_{B_v}$. Regularize first and use Tonelli. The comparison-field masses are

$$\mu_Y(B_v) = C^{-g/2} r_v^{d+g/2} e^{\gamma G_v} Z_v,$$

where the independent Z_v have all negative moments. Each tree increment appears with cumulative exponent P_v. Taking Gaussian expectations in (5.7) gives the factor

$$(r_v / r_{\text{par}(v)})^{P_v(gc_v - d - g/2) - gP_v^2/2} = (r_v / r_{\text{par}(v)})^{\frac{g}{2}(c_v - c)_+^2}$$

at each proper node.

For a cluster of m charged leaves, the pair density, the configuration volume, and this inverse-mass gain combine to

$$E_m = d(m-1) - \frac{g}{2}m(m-1) + \frac{g}{2}(m-c)_+^2. \tag{5.8}$$

If $m \le c$, then $E_m = (m-1)(d - gm/2) > 0$. If $m > c$, direct expansion gives the dimension-independent-in-m positive value

$$E_m = \frac{(2d-g)^2}{8g} > 0. \tag{5.9}$$

For a cluster containing the uncharged root and j charged leaves, the corresponding exponent is

$$E_j^0 = dj - \frac{g}{2}j(j-1) + \frac{g}{2}(j-c)_+^2 > 0;$$

above threshold it equals $gc^2/2$.

For completeness, the volume factor uses representatives inherited from one child at each node, with root representative $z_0 = 0$. A node with ℓ_v children introduces $\ell_v - 1$ independent d-dimensional coordinate increments of size $O(r_v)$. The linear change of coordinates has Jacobian one. The product volume bound $C\prod_v r_v^{d(\ell_v-1)}$, rewritten in scale ratios, contributes $d(m_v-1)$ at every proper node. There are finitely many tree shapes and labels. Quantizing ratios dyadically and using the positive exponents proves (5.6).

5.3 The radial comparison and exact Gaussian cancellation

Split ordered k-tuples of potential neighbors according to

$$\max_i |q_i - p| \in (r, 2r], \qquad r = \rho 2^{-j-1},$$

and write $q_i = p + rz_i$, $T = \log(1/r)$, and

$$a_e = d + g/2 - eg > 0.$$

The multipoint Palm formula gives the radial prefactor

$$t^k r^{ka_e - gk^2/2}. \tag{5.10}$$

The lower background cloud avoids the union of screening balls. Those balls avoid every inserted point, so restrictions caused by inserted points with marks at most s may be dropped for an upper bound. The remaining void factor is the exponential of minus s times the inserted GMC mass. Bounded covariance remainders only introduce constants depending on k.

The rescaled original covariance is $T + \log|x-y|^{-1} + R(p+rx, p+ry)$. Choose a fixed C so that the covariance of

$$Y_r = \Omega + F, \qquad \Omega \sim N(0, T+C), \quad \Omega \perp F,$$

dominates it, uniformly in p, r. Applying Kahane comparison to the convex void functional bounds the relevant expectation above by

$$\mathbb{E}\exp[-c_k s r^{a_e - gk} e^{\gamma\Omega} M_z]. \tag{5.11}$$

This step does not identify Ω with an actual sphere average of h.

For $\sigma^2 = T + C$, Gaussian tilting gives exactly

$$\begin{aligned}&\mathbb{E}\exp[-c_k s r^{a_e - gk} e^{\gamma\Omega} M] \\ &\quad = \mathbb{E}\left[e^{k\gamma\Omega - gk^2(T+C)/2} \exp[-c_k e^{-gkC} s r^{a_e} e^{\gamma\Omega} M]\right].\end{aligned} \tag{5.12}$$

Indeed, the tilt translates Ω by $k\gamma(T + C)$. Multiplying (5.12) by (5.10) cancels the apparently divergent factor $r^{-gk^2/2}$. The radial factorial moment is therefore bounded by

$$C_k(t/s)^k \mathbb{E} H_k(\log s - a_e T + \gamma\Omega + c_k'), \tag{5.13}$$

where

$$H_k(u) = e^{ku} \int W(z) \mathbb{E}[e^{-e^u M_z}]\, dz.$$

By Tonelli and (5.6),

$$\int_{\mathbb{R}} H_k(u)\, du = \Gamma(k) \int W(z) \mathbb{E}[M_z^{-k}]\, dz < \infty. \tag{5.14}$$

5.4 Summing all scales

The arguments in (5.13) have the marginal distributions of a Brownian motion with negative drift, sampled on a fixed time lattice. For $a > 0$, the density

$$f_u(y) = \frac{1}{\sqrt{2\pi g u}} \exp\left[-\frac{(y + au)^2}{2gu}\right]$$

satisfies

$$\int_0^\infty f_u(y)\, du = \frac{1}{a} e^{-a(y+|y|)/g} \le \frac{1}{a}.$$

As a function of u, it is unimodal. Its supremum for $u \ge u_0 > 0$ is uniformly bounded in y. Hence its sum over any fixed lattice starting at u_0 is uniformly bounded by a constant times this integral plus the supremum. No independence between different radial scales is used.

Apply this bound to (5.14). Summing (5.13) proves

$$\mathbb{E}[(L_{s,t}(p))_k] \le C_k (t/s)^k$$

for every integer factorial moment. Expand ordinary integer powers into factorial powers and use $t/s \ge 1$. Jensen's inequality applied to the next integer moment gives (5.1) for every real $q \ge 1$. Mecke gives (5.2). This proves Theorem 5.1.

Remark 5.2 (Scope of the moment estimate). The count being controlled is a local set of sites with empty-ball witnesses. It is not the total number of Poisson points in a fixed domain. The latter inherits the finite positive-moment range of GMC and need not have moments of every order. This distinction is essential.

6 Mass-ball comparisons and endpoint repair

Let $L_\delta(x, y; U)$ be the minimum number of Euclidean balls in U, each of μ-mass at most δ, whose union covers a continuous path from x to y. Let $\overline{L}_\delta$ denote the variant requiring $\mu(2B) \le \delta$ for each ball.

Proposition 6.1 (Subpower comparison)**.** *For $U \Subset V$ and every $\varepsilon > 0$, eventually almost surely in the marked coupling,*

$$d_\lambda(x, y; V) \le \lambda^{\varepsilon}\big(\overline{L}_{1/\lambda}(x, y; U) + 1\big), \qquad x, y \in U. \tag{6.1}$$

For graph paths with nuclei in $K \Subset U$, the reverse comparison is

$$L_{C \log \lambda/\lambda}(x, y; U) \le d_\lambda(x, y; K) + 2. \tag{6.2}$$

The notation on the right of (6.2) restricts graph-path nuclei to K; endpoint conventions only affect the displayed bounded error.

Proof. Each Delaunay edge lies in a closed empty witness ball. Two further nearest-site empty balls cover the endpoint-to-nucleus segments. Lemmas 4.1 and 4.2 keep these balls inside U and give (6.2). Sphere-nullity permits arbitrarily small enlargements if open covering balls are required.

For the upper bound, first observe that the number of Voronoi cells meeting $B_r(q)$ is at most

$$\#(\mathcal{P} \cap B_{2r}(q)) + \deg^{\mathrm{weak}}_{\mathcal{P}\cup\{q\}}(q). \tag{6.3}$$

If a far nucleus $p \notin B_{2r}(q)$ has a cell meeting the ball at y, the empty ball centered at y with p on its boundary contains q in its interior. Shrink it by a homothety about p until q is on its boundary. The shrunken empty ball witnesses the second term in (6.3).

On an intensity interval $[L, 2L]$, choose a deterministic lower radius $r_0 = L^{-M}$ with M large enough that the uniform upper mass bound makes every required doubled r_0-ball have mass at most $(2L)^{-1}$. Enlarge smaller admissible balls to this scale. Larger balls can be covered by a bounded number of smaller balls with centers and radii rounded to a deterministic polynomial-size grid, keeping their doubled balls inside the original doubled ball. Conditional Poisson tails bound all relevant doubled-ball counts by $O(\log L)$. Theorem 5.1, applied with lower intensity L and upper intensity $2L$, and a union bound over the deterministic grid, bound all relevant weak query degrees by L^u for any prescribed $u > 0$. Choose the moment order large enough to make the dyadic failure probabilities summable. Uniform filling localizes the empty-ball witnesses to the fixed buffer.

Thus each admissible mass ball meets at most L^{ε} cells after taking a slightly smaller intermediate exponent. A connected curve in the union of N such balls meets a connected collection of at most NL^{ε} cells. In general position, cells sharing a lower-dimensional face correspond to a Delaunay simplex, whose vertices are pairwise adjacent. The cell collection therefore yields a graph path of the claimed cost. Uniform mesh gives confinement in V. Dyadic Borel–Cantelli proves the assertion for all sufficiently large real intensities. □

6.1 Temporary guards

Inserting a site can add at most one cell to the set of cells covering a fixed trace: every other new cell is contained in its old counterpart. Deleting a guard only changes the partition inside its former cell. Every new nucleus whose cell occupies part of that region admits a weak empty-ball witness with the deleted guard. Thus the trace can be repaired with an overhead bounded by the corresponding weak-neighbor count, plus the finitely many remaining guards. Deleting a finite guard set G costs at most

$$\sum_{p \in G} L^0_{s,t}(p) + C|G|^2$$

in a suitably screened buffer. Theorem 5.1 supplies moments for this overhead. The argument uses convex cells and empty balls, not planar intersections.

6.2 The final attachment scale

The same proof of (6.3), with the uniform upper GMC power, gives the following useful endpoint conclusion. For any $u > 0$, choose M large enough and let $r_{\min} = \lambda^{-M}$. With probability tending to one, uniformly in centers in a fixed compact, the representatives of any two query points in a radius-$Cr_{\min}$ ball can be joined with at most λ^u edges. One first uses a continuous trace inside that small ball; its complete family of intersecting cells is facet-connected. The graph route lies in the enlargement of the ball by the largest Euclidean cell diameter on the fixed buffer, which tends to zero. It need not lie in a fixed multiple of the microscopic ball itself: its nearest nuclei can be farther away than $r_{\min}$. Round ball centers to a deterministic grid and use all-order screening exactly as above. The anchor construction below keeps this vanishing mesh enlargement when establishing confinement in a prescribed open chart.

Consequently, if the chosen normalizing cost satisfies $B_\lambda \geq \lambda^\alpha$ eventually for some $\alpha > 0$, taking $u < \alpha$ makes these final attachment costs negligible in B_λ units. What is still missing is the cost control at all larger intermediate scales.

7 Coarse Gaussian profiles: an obstruction and a multiscale repair

The uniform L^2 covariance bound for the coarse gradient used in the planar argument fails in dimensions $d > 2$. A different estimate nevertheless gives exponential tails for the diameter of clusters of uncontrolled coarse blocks. This section proves that Gaussian estimate. Applying it to the Delaunay metric still requires a local connection event whose conditional failure probability can be made small.

7.1 The Gaussian decomposition and the failure of the gradient bound

Let $\kappa \in C_c^\infty(\mathbb{R}^d)$ be real, with $\int \kappa \neq 0$ and $\int \kappa^2 = 1$, and set $k = \kappa * \widetilde{\kappa}$, where $\widetilde{\kappa}(x) = \kappa(-x)$. All constants below may depend on d and κ. Let W be real white noise on $\mathbb{R}^d \times (0, \infty)$ with control measure $dy\, ds$. For $L \geq 2$, define the smooth coarse field

$$H_L(u) = \int_1^L \int_{\mathbb{R}^d} s^{-(d+1)/2} \kappa\Big(\frac{u-y}{s}\Big)\, W(dy, ds). \tag{7.1}$$

Its covariance is

$$C_L(u-v) = \int_1^L k\Big(\frac{u-v}{s}\Big)\, \frac{ds}{s}.$$

We use the Fourier convention $\widehat{f}(\zeta) = \int e^{-iu\cdot\zeta} f(u)\, du$. The covariance spectral density is

$$f_L(\zeta) = \int_1^L s^{d-1} |\widehat{\kappa}(s\zeta)|^2\, ds. \tag{7.2}$$

Proposition 7.1 (Failure of the planar gradient estimate)**.** *If $d > 2$, the $L^2(\mathbb{R}^d; \mathbb{R}^d)$ operator norm of the covariance of ∇H_L is comparable to L^{d-2}. Moreover, there are cubes A_L of side comparable to L such that, for every fixed $\theta > 0$ and all sufficiently large L,*

$$\mathbb{E} \exp\Big(\theta \int_{A_L} |\nabla H_L(u)|^2\, du\Big) = \infty. \tag{7.3}$$

In particular, a uniform positive exponential-moment parameter for this quadratic integral is unavailable in dimensions $d > 2$.

Proof. The covariance multiplier of the gradient is $\zeta\zeta^{\mathsf{T}} f_L(\zeta)$, so its norm is $\sup_\zeta |\zeta|^2 f_L(\zeta)$. Continuity and $\widehat{\kappa}(0) \neq 0$ imply that, for a sufficiently small fixed $\eta > 0$, $|\widehat{\kappa}(s\eta e_1/L)|$ is bounded below for $1 \le s \le L$. At $\zeta = \eta e_1/L$ this gives the lower bound cL^{d-2}. For $|\zeta| \le L^{-1}$, boundedness of $\widehat{\kappa}$ gives the matching upper bound. For $|\zeta| \ge L^{-1}$, rapid Fourier decay gives, uniformly in the direction of ζ,

$$f_L(\zeta) \le \int_0^\infty s^{d-1} |\widehat{\kappa}(s\zeta)|^2 \, ds \le C|\zeta|^{-d}.$$

Since $d > 2$, $|\zeta|^{2-d} \le L^{d-2}$ in this range.

For the assertion on a finite region, fix a nonzero real $\phi \in C_c^\infty(\mathbb{R}^d)$ with $\|\phi\|_2 = 1$, and put $v_L(u) = L^{-d/2}\phi(u/L)e_1$. Choose A_L containing its support. Parseval's identity and the substitutions $\eta = L\zeta$ and $t = s/L$ give

$$\begin{aligned} \operatorname{Var}\langle \nabla H_L, v_L \rangle &= (2\pi)^{-d} L^{d-2} \int_{\mathbb{R}^d} |\widehat{\phi}(\eta)|^2 \eta_1^2 \left[\int_{1/L}^1 t^{d-1} |\widehat{\kappa}(t\eta)|^2 \, dt \right] d\eta \\ &\ge c_\phi L^{d-2} \end{aligned}$$

for all sufficiently large L. The limiting integral in brackets produces a strictly positive integral: $\widehat{\kappa}$ is nonzero near the origin, and a nonzero compactly supported smooth ϕ cannot have $\eta_1 \widehat{\phi}(\eta) = 0$ almost everywhere. By Cauchy–Schwarz, $\int_{A_L} |\nabla H_L|^2 \ge |\langle \nabla H_L, v_L \rangle|^2$. For a centered Gaussian variable of variance σ^2, $\mathbb{E}e^{\theta Z^2} = \infty$ when $2\theta\sigma^2 \ge 1$. This proves (7.3). □

The multiplier $|\zeta|^d f_L(\zeta)$ is uniformly bounded. This critical-order observation does not by itself control the local profile needed for a change of measure. Multiplication by a cutoff creates lower-order derivative terms; subtracting a block constant leaves affine and higher polynomial profiles. The argument below controls all the required derivatives locally instead.

7.2 The local Cameron–Martin norm

Let X be the independent fine reference field obtained by integrating the same white-noise kernel over $0 < s < 1$, in the sense of distributions. Its spectral density is

$$f_X(\zeta) = \int_0^1 s^{d-1} |\widehat{\kappa}(s\zeta)|^2 \, ds.$$

The assumption $\widehat{\kappa}(0) \neq 0$ and rapid Fourier decay imply

$$c(1 + |\zeta|^2)^{-d/2} \le f_X(\zeta) \le C(1 + |\zeta|^2)^{-d/2}. \tag{7.4}$$

Indeed, for $|\zeta| \ge 1$ the upper bound follows by the substitution $t = s|\zeta|$, while integration over $0 < s < c_0/|\zeta|$ gives the lower bound. On $|\zeta| \le 1$ one uses a fixed sufficiently small interval of s. Consequently, for smooth compactly supported f,

$$\|f\|^2_{\mathrm{CM}(X)} = (2\pi)^{-d} \int \frac{|\widehat{f}(\zeta)|^2}{f_X(\zeta)} \, d\zeta \asymp \|f\|^2_{H^{d/2}(\mathbb{R}^d)}. \tag{7.5}$$

Fix $b \ge 1$, a smooth cutoff χ supported in $(-b, b)^d$, and an inner region on which $\chi = 1$. For $i \in \mathbb{Z}^d$ write $B_i = i + [-b, b]^d$, $\chi_i(u) = \chi(u - i)$, and

$$\overline{H}_{L,i} = |B_i|^{-1} \int_{B_i} H_L(u) \, du, \qquad g_{L,i} = \chi_i (H_L - \overline{H}_{L,i}).$$

With $m = \lceil d/2 \rceil$, the product rule, the embedding $H^m \subseteq H^{d/2}$, and the mean value theorem give

$$\|g_{L,i}\|_{\mathrm{CM}(X)} \le C_{\mathrm{loc}} \max_{1 \le |\alpha| \le m} \sup_{u \in B_i} |\partial^\alpha H_L(u)|. \tag{7.6}$$

Here and below α ranges over multi-indices. The zeroth-order term in the product rule is bounded using $\sup_{B_i} |H_L - \overline{H}_{L,i}| \le \operatorname{diam}(B_i) \sup_{B_i} |\nabla H_L|$. Thus (7.6) is valid in odd dimensions as well.

7.3 Rare obstacles associated with independent scale bands

Set $R_j = 2^j$, $j \ge 0$, and split (7.1) into the independent fields

$$H_{L,j}(u) = \int_{R_j}^{\min(2R_j, L)} \int s^{-(d+1)/2} \kappa\Big(\frac{u-y}{s}\Big) W(dy, ds), \qquad R_j < L. \tag{7.7}$$

Time intervals may be taken half open; endpoints have zero white-noise control measure. In particular all the scale bands are independent, and $H_L = \sum_{R_j < L} H_{L,j}$.

Use the notation $Q(z,r) = z + [-r,r]^d$. Choose fixed constants $a > b + 1$ and $A > a + 2$. For $z \in R_j \mathbb{Z}^d$ define

$$Z_{L,j,z} = \max_{1 \le |\alpha| \le m} R_j^{|\alpha|} \sup_{u \in Q(z, aR_j)} |\partial^\alpha H_{L,j}(u)|.$$

A code (j,z) is *active* if

$$Z_{L,j,z} > M\sqrt{R_j}. \tag{7.8}$$

To an active code attach the closed obstacle $Q(z, AR_j)$, and write $\mathcal{O}_{L,M}$ for their union.

Lemma 7.2 (Activation tails and dependence)**.** *There exist $M_0, c_0 > 0$ and an integer $C_0 \ge 1$, depending only on the fixed data, such that for $M \ge M_0$,*

$$\mathbb{P}((j,z) \text{ is active}) \le \exp(-c_0 M^2 R_j). \tag{7.9}$$

At each fixed scale the grid codes admit a coloring with C_0 colors such that the activation events within each color are mutually independent. The collections at distinct scales are independent. For any prescribed finite set of distinct codes $\mathcal{T}$,

$$\mathbb{P}(\mathcal{T} \text{ is entirely active}) \le \exp\left(-c_1 M^2 \sum_{(j,z) \in \mathcal{T}} R_j\right), \qquad c_1 = c_0/C_0. \tag{7.10}$$

All constants are uniform in L, including its last truncated band.

Proof. The spatial and scale substitutions $u = R_j v$, $y = R_j w$, $s = R_j t$ show that the scaled derivatives appearing in $Z_{L,j,z}$ have the law of the derivatives of the band with scales $1 \le t \le b_j$, where $1 < b_j \le 2$. For a complete band $b_j = 2$ the law is independent of j and z. For all truncated bands, the pointwise variances of every derivative are uniformly bounded. Applying a deterministic Sobolev embedding on a slightly enlarged fixed cube to finitely many derivatives, and then Cauchy–Schwarz, gives $\mathbb{E} Z_{L,j,z} \le C$ uniformly in $b_j \in (1,2]$. The pointwise variances of the centered Gaussian process whose supremum defines $Z_{L,j,z}$ are also uniformly bounded. Gaussian supremum concentration therefore gives

$$\mathbb{P}(Z_{L,j,z} > \mathbb{E} Z_{L,j,z} + t) \le e^{-t^2/(2\sigma^2)}.$$

Increase M_0 so that $M\sqrt{R_j} - \mathbb{E}Z_{L,j,z} \geq M\sqrt{R_j}/2$. This proves (7.9).

Choose r_κ with $\operatorname{supp}\kappa \subseteq [-r_\kappa, r_\kappa]^d$. The activation event uses only white noise in

$$Q(z, (a+2r_\kappa)R_j) \times [R_j, \min(2R_j, L)].$$

Color $z/R_j \in \mathbb{Z}^d$ by its coordinatewise residues modulo an integer $K > 2(a+2r_\kappa)$. Within a color these spatial supports are disjoint. One may take $C_0 = K^d$. For a prescribed set having n_j codes at scale j, one color contains at least n_j/C_0 of them. The probability that all codes at this scale are active is at most $e^{-c_0 M^2 R_j n_j/C_0}$ by independence in that color. Independence of the time bands now gives (7.10). □

Lemma 7.3 (Obstacles cover uncontrolled local profiles)**.** *There is a fixed $C_* > 0$ such that each lattice unit cube $Q(i, 1/2)$ with*

$$\|g_{L,i}\|_{\mathrm{CM}(X)} > C_* M \tag{7.11}$$

is contained in an active obstacle. In particular connected clusters of such unit cubes are contained in connected components of $\mathcal{O}_{L,M}$.

Proof. For each scale choose a nearest grid center $z_j(i) \in R_j\mathbb{Z}^d$, with deterministic tie breaking, so that $|z_j(i) - i|_\infty \leq R_j/2$. Our choice of a ensures $B_i \subseteq Q(z_j(i), aR_j)$. If all these codes are inactive, then for $1 \leq |\alpha| \leq m$,

$$\sup_{B_i} |\partial^\alpha H_L| \leq M \sum_{j \geq 0} R_j^{1/2-|\alpha|} \leq \frac{M}{1 - 2^{-1/2}}.$$

Together with (7.6), this contradicts (7.11) when $C_* = C_{\mathrm{loc}}/(1-2^{-1/2})$ is increased to accommodate the fixed multi-index norm convention. Hence some code $(j, z_j(i))$ is active. Because $R_j \geq 1$ and $A > 2$, its obstacle contains $Q(i, 1/2)$. Two adjacent closed unit cubes intersect, so their containing obstacles lie in the same component. □

7.4 An exponential diameter estimate

Theorem 7.4 (Uniform coarse-obstacle diameter tails)**.** *There are constants $M_1, C, c > 0$ such that, for all $M \geq M_1$, $L \geq 2$, $i \in \mathbb{Z}^d$, and $N \geq 2$,*

$$\begin{aligned}\mathbb{P}\big(&\text{some component of } \mathcal{O}_{L,M} \text{ meets } Q(i, 1/2)\\ &\text{and meets } \mathbb{R}^d \setminus Q(i, N)\big) \leq Ce^{-cM^2N}.\end{aligned} \tag{7.12}$$

The same estimate, with changed constants, holds for the diameter of the cluster of unit cubes satisfying (7.11) and containing a specified unit cube.

Proof. For clarity translate the specified cube to $Q(0, 1/2)$. A chain is an ordered list of obstacles with consecutive obstacles intersecting. A chain may always be reduced to a simple chain, with distinct grid/scale codes. For a prescribed simple chain with radii $R_1, \ldots, R_n$, (7.10) bounds its activation probability by $\exp(-c_1 M^2 \sum_{\ell=1}^n R_\ell)$.

For a given first radius $R_1 \geq 1$, at most a fixed number of grid centers have their obstacle intersect $Q(0, 1/2)$. Given an obstacle center of radius R, the number of possible centers on the next grid of spacing R' whose obstacles intersect it is at most

$$C_A(1 + R/R')^d \leq C_A e^{dR}, \qquad R, R' \geq 1. \tag{7.13}$$

Consequently the total number of center sequences for a prescribed radius sequence is at most $\exp(B\sum_\ell R_\ell)$ for a fixed B. This bound may overcount nonsimple center sequences; probability estimates are applied only to simple chains before this numerical overcounting.

If a chain starts at $Q(0,1/2)$ and reaches outside $Q(0,N)$, the triangle inequality along the intersecting cubes gives

$$2A\sum_{\ell=1}^{n} R_\ell \ge N - 1/2, \qquad \text{hence} \qquad \sum_{\ell=1}^{n} R_\ell \ge c_A N \quad (N \ge 2).$$

Set $q = c_1 M^2 - B$. Union bound, reserve half the exponential weight for the preceding geometric constraint, and then sum the radii:

$$\begin{aligned}\mathbb{P}(\text{such an active chain exists}) &\le \sum_{n\ge 1}\sum_{R_1,\dots,R_n\in\{1,2,4,\dots\}} \mathbf{1}_{\{\sum R_\ell \ge c_A N\}} e^{-q\sum R_\ell} \\ &\le e^{-qc_A N/2}\sum_{n\ge 1}\left(\sum_{j\ge 0} e^{-q2^j/2}\right)^n .\end{aligned}$$

Choose M_1 so that $q \ge c_1 M^2/2$ and the inner sum is at most $1/2$. The geometric series is then bounded uniformly, proving (7.12).

There is no issue concerning accumulation of obstacles. For every fixed compact set, the expected number of active obstacles intersecting it is finite, since the number of possible centers at scale R_j is at most $C(1+\operatorname{diam}(K)/R_j)^d$ and (7.9) is summable over j. Thus the family is locally finite almost surely, even if all scales are allowed. Connected components are therefore represented by the intersection graph of obstacles, and any connection in question has a finite chain of intersecting obstacles. Finally apply Lemma 7.3. □

7.5 Adding conditional fine-scale failures

The same argument permits local failures of the fine-scale construction on the controlled blocks. This extension separates the Gaussian issue from the as-yet-required metric connection theorem.

Corollary 7.5 (Mixed coarse and fine bad blocks)**.** *Let $\mathcal{G}$ be the sigma field of the entire coarse noise. Suppose F_i, $i \in \mathbb{Z}^d$, are fine-failure events, defined to be false on uncontrolled blocks, and suppose there is a fixed coloring with C_f colors such that, conditionally on $\mathcal{G}$, the events of each color are mutually independent and*

$$\mathbb{P}(F_i \mid \mathcal{G}) \le p \qquad \textit{almost surely.} \tag{7.14}$$

Add $Q(i,A)$ as a radius-one obstacle whenever F_i occurs. If M is sufficiently large and p sufficiently small, the resulting obstacle union has a diameter bound Ce^{-cN} as in Theorem 7.4, uniformly in L.

Proof. For prescribed distinct fine sites S, one color contains at least $|S|/C_f$ sites, so

$$\mathbb{P}\left(\bigcap_{i\in S} F_i \mid \mathcal{G}\right) \le p^{|S|/C_f}.$$

For any prescribed collection $\mathcal{T}$ of coarse codes, multiply this conditional bound by the indicator that all its codes are active and average. Lemma 7.2 gives the joint estimate

$$\mathbb{P}\left(\bigcap_{i\in S} F_i,\ \mathcal{T} \text{ entirely active}\right) \le \exp\left[-c_1 M^2 \sum_{(j,z)\in\mathcal{T}} R_j - \frac{\log(1/p)}{C_f}|S|\right].$$

Use simple chains of distinct typed codes. Fine obstacles have radius one, with one additional possible type at that radius. The center bound (7.13) remains valid and this extra type contributes only a factor at most 2^n to a length-n chain. The preceding proof applies once $\min\{c_1 M^2, C_f^{-1}\log(1/p)\}$ dominates its fixed counting constant. □

For completeness, the role of the controlled Cameron–Martin norm is quantitative. If E is an event for the fine field, then for a deterministic shift $g \in \mathrm{CM}(X)$ the Gaussian density formula and Cauchy–Schwarz give

$$\mathbb{P}(X + g \in E) \le e^{\|g\|^2_{\mathrm{CM}(X)}/2}\, \mathbb{P}(X \in E)^{1/2}. \tag{7.15}$$

On controlled blocks the prefactor is at most $e^{C_*^2 M^2/2}$. Thus one first fixes a sufficiently large M for the coarse-obstacle estimate, and then requires a sufficiently small unshifted local failure probability to obtain the p in (7.14). The same calculation works conditionally on independent auxiliary randomness. Conditional finite-range independence of local fine events must come from the fine-field/noise construction; it is not inferred from the Cameron–Martin density.

Remark 7.6 (Scope of the repair). The result replaces the unavailable gradient quadratic-moment bound by a direct cluster-diameter estimate. It does not assert the planar volume-exponential estimate for arbitrary prescribed sets of coarse bad blocks. It also does not supply a connected crossing or port network for the Delaunay hop metric in dimensions $d \ge 3$. Such a metric event, with a sufficiently small failure probability and the necessary endpoint attachment, is a separate hypothesis in Corollary 7.5. Accordingly this section repairs an auxiliary Gaussian/percolation input and does not establish metric compactness or convergence.

8 A dimension-free replacement for planar joining

Lemma 8.1 (Dyadic anchor joining). *Let $r_j = 2^{-j}$, $j_0 \le j \le J$. At each dyadic grid point z in a fixed buffered region, use the same graph representative $v_\lambda(z)$. Suppose that:*

1. *neighboring level-j anchors, and every level-$(j+1)$ anchor and its level-j parent, have graph connections of cost at most $B_\lambda A r_j^\theta$, confined within Euclidean distance $C_d r_j$ of the anchors;*
2. *every query representative is connected to its level-J grid anchor at cost at most $B_\lambda e_\lambda$, within distance $C_d r_J$.*

Here $\theta > 0$ and the neighbor list contains the finite coordinate moves needed to join grid cubes at distance at most $2r_j$. Then, for points in the inner region with $|x-y| \le r_{j_0}/C_d$,

$$B_\lambda^{-1} d_\lambda(x,y) \le C_{d,\theta} A(|x-y| + r_J)^\theta + 2e_\lambda. \tag{8.1}$$

The constructed path stays in a $C_d(|x-y| + r_J)$ neighborhood of the endpoints. In particular it is internal to any prescribed open chart if the initial scale is smaller than its clearance.

Proof. If $|x-y| \ge r_J$, choose a level with r_j comparable to $|x-y|$. Their grid anchors differ by only a bounded number of neighboring coordinate moves. Join these anchors, then join each to its finer anchors successively and finally to the query representative. The total cost is bounded by

$$C_d B_\lambda A \sum_{k=j}^{J} r_k^\theta + 2B_\lambda e_\lambda \le C_{d,\theta} B_\lambda A r_j^\theta + 2B_\lambda e_\lambda.$$

If $|x-y| < r_J$, use neighboring finest-grid anchors. The same geometric-series bound proves confinement. No common ancestor assumption about nearby points in a fixed dyadic tree is needed: the neighboring-anchor move at the starting scale handles grid boundaries. □

This lemma constructs actual intersections by using exactly the same graph vertex at every shared port. Temporary deterministic guards are an alternative way to impose common ports; their subsequent deletion is controlled by screening. Existence of unrelated cheap crossings is insufficient to verify its hypotheses.

The next section proves the required local joining estimates by high-quantile percolation, using the lattice geometry of [7, Section 4.1]. Adjacent scales share actual deterministic ports, and a past-intensity bound controls the normalizers.

9 A past-envelope estimate for finite anchor networks

Here is a more concrete replacement for the planar crossing amplification. It uses finite deterministic ports, the coarse-obstacle estimate, and the adjacent-scale open-cluster construction of [7, Section 4.1]. It does not assume point-to-point moment bounds at the crossing normalization. The normalizer is instead defined from the same finite port observable used to create the open clusters. The screening estimate is proved above, and the microscopic attachment estimates are derived below from it and standard GMC moments. The tail estimate at a general intensity has an explicit past-envelope hypothesis; the final corollary supplies deterministic record sequences satisfying that hypothesis. These results do not identify the normalizing exponent with that of the target metric.

9.1 The observable and the geometric hypotheses

Fix $d \geq 2$ and $0 < \gamma < \sqrt{2d}$, and write

$$b = d + \gamma^2/2, \qquad Q = b/\gamma, \qquad \beta_\eta = b - \gamma(\sqrt{2d} + \eta) > 0$$

for a fixed $0 < \eta < Q - \sqrt{2d}$. Let X be the stationary compact-kernel field with noise scales in $(0,1)$, normalized so that its logarithmic covariance coefficient is one. Its chaos has the exact decomposition, at a cutoff r,

$$\mu_X(z + r\,du) = r^b e^{\gamma H_r(z+ru)} \mu_{X^{\mathrm{fine}}}(du),$$

where the rescaled fine field has the law of X and is independent of the noise at scales at least r.

Here is one explicit finite port scheme. Write $B_\infty(z,r) = z + (-r,r)^d$ and choose a sufficiently large fixed integer m_0. For $\theta \in \{0,1\}$ set

$$P_\theta = 2^{-\theta}\mathbb{Z}^d \cap [-4,4]^d, \qquad G_\theta = (2^{-\theta}/m_0)\mathbb{Z}^d \cap [-16,16]^d.$$

Include every oriented pair (p,q) in P_θ with $q - p = \pm 2^{-\theta} e_k$, and confine its trace to $B_\infty(p, 2^{1-\theta})$. Test that trace in the diagram with guard set G_θ. The observation box can be enlarged once, if needed for the support radius of the fixed field kernel. Taking m_0 large makes all cells meeting the trace domains depend only on sites in this fixed observation box: the dense guards provide a closer site than any exterior point. The local diagrams agree with the corresponding infinite periodic guard pattern in their trace domains.

Translations by the relevant lattice vertices preserve these guard patterns. When adjacent lattice scales are inspected using the coarser cutoff, the coarser tests use $\theta = 0$ and the surrounding

finer tests use $\theta = 1$. All their endpoints and trace domains occur in the displayed finite list. Thus no arbitrary-endpoint distance is included in this definition.

For each encoded connection, its cell-trace cost means the minimum number of *distinct closed Voronoi cells* whose union contains a continuous connection between its ports inside the prescribed open trace cube. We use this cover convention throughout; in particular we do not count raw cell transitions along a trace which can reenter a cell. Let $F_I \geq 1$ be one plus the sum of these finitely many costs for a cloud of intensity $I\mu_X$. Different encoded guard patterns may be tested on the same Poisson cloud. A straight trace and the finite observation buffer give

$$\mathbb{E}F_I \leq C(1 + I). \tag{9.1}$$

Every local connection used below is dominated by the corresponding translated and rescaled copy of this finite observable. Guards are only an auxiliary device: traces from neighboring tests have the same deterministic port as an endpoint, and deleting guards does not change the geometric trace.

The following are the precise microscopic estimates used in the argument. They concern the original Poisson geometry and do not assert a normalized graph-distance modulus.

1. For every $B, u > 0$, a sufficiently small polynomial radius can be chosen so that, with probability $1 - O(I^{-B})$, every ball of that radius times a fixed power of $\log I$, in the construction buffer, meets at most I^u cells. The assertion includes the finitely many permanent macro guards.

2. For a fixed finite collection of dyadic guard grids down to level $J \leq M \log_2 I + O(1)$, all auxiliary guard deletion costs are at most I^u each with probability $1 - O(I^{-B})$, for any fixed $B, u, M > 0$. One may use the weaker total bound $gI^u + g^2$ for a construction involving at most g guards.

Lemma 9.1 (Quantitative terminal mesh and guard repair)**.** *The two preceding estimates follow from the uniform all-order rootless screening bound and the usual positive and negative moments of subcritical GMC on fixed compact buffers. The finitely many macro guards can also be removed at cost $O(I^u)$, with failure $O(I^{-B})$ for every prescribed $u, B > 0$. Furthermore, there exists $c_0 > 0$ such that*

$$\mathbb{P}(F_I \geq I^{c_0}) \longrightarrow 1. \tag{9.2}$$

Proof. Write

$$\zeta(q) = (d + \gamma^2/2)q - \gamma^2 q^2/2.$$

Uniformly over a fixed compact buffer, $\mathbb{E}\mu_X(B_r(z))^q \leq C_q r^{\zeta(q)}$ for $1 < q < 2d/\gamma^2$. Since $\zeta(1) = d$ and $\zeta'(1) = d - \gamma^2/2 > 0$, choose such a q with $\zeta(q) > d$. Take $r = I^{-M}(\log I)^c$, where c is any fixed exponent needed for the terminal cluster displacement. A deterministic r-net has $O(r^{-d})$ points. Markov's inequality and a union bound give

$$\mathbb{P}\left[\max_{z \text{ in the net}} \mu_X(B_{Cr}(z)) > I^{-1}\right] \leq CI^q r^{\zeta(q)-d}.$$

For any desired B, choose M sufficiently large to make this $O(I^{-B})$, absorbing the fixed logarithmic power. Conditional on the field and this mass event, the Poisson counts in the enlarged net balls have bounded means. A Poisson tail and the same polynomial-size union bound make all those counts at most $I^{u/2}$ with failure smaller than any prescribed power of I^{-1}.

For a deterministic query z, the elementary cell-count bound is

$$\#\{\text{cells meeting } B_r(z)\} \le \#(\mathcal{P}_I \cap B_{2r}(z)) + L_I(z).$$

Here $L_I(z)$ is the rootless potential-neighbor count after inserting z; one may use the ordinary inserted-query degree when no localization is needed. The inequality follows by shrinking an empty ball centered at a cell contact towards a nucleus outside $B_{2r}(z)$ until its boundary first reaches z. Fixed-radius filling localizes the needed empty-ball witnesses. All their moments are uniformly bounded, so a sufficiently large moment order and the polynomial net cardinality make $\max_z L_I(z) \le I^{u/2}$ with failure $O(I^{-B})$. Rounding an arbitrary center to this net and enlarging the ball by a fixed factor proves the uniform cell-count assertion. Inserting a fixed number of macro guards adds at most that number of cells.

For auxiliary guards, the union of all candidate grids through level $J \le M \log_2 I + O(1)$ has at most CI^{dM} sites. Keeping a fixed dense macro guard net makes every relevant empty-ball witness local deterministically. Another union bound in an arbitrarily high rootless moment gives the asserted deletion event. To remove the macro guards at the end, a finite fixed-radius filling cover fails with probability at most

$$\sum_B \mathbb{E}e^{-I\mu_X(B)} \le C_k I^{-k} \sum_B \mathbb{E}\mu_X(B)^{-k}$$

for every $k > 0$. On its complement, the same local empty-ball witnesses control their deletion, proving the final-guard assertion.

Finally, the negative moment bound at shrinking radii is

$$\mathbb{E}\mu_X(B_r(z))^{-k} \le C_k r^{-(d+\gamma^2/2)k-\gamma^2k^2/2}.$$

For $k = 1$, the probability that some ball in an r-net is empty is at most $CI^{-1}r^{-(2d+\gamma^2)}$. Choose $r = I^{-a}$ with $0 < a < 1/(2d+\gamma^2)$. With probability tending to one, all these balls contain a Poisson site. Hence every cell meeting the smaller construction buffer has diameter at most Cr, including after insertion of fixed guards. A trace joining one fixed separated port pair needs at least c/r cells. Taking any $0 < c_0 < a$ proves (9.2). □

Theorem 9.2 (Anchor-network tails from a past envelope)**.** *Assume the all-order rootless screening bound used in Lemma 9.1 and the coarse-profile obstacle estimate of Section 7. Let $v(J) \ge 1$ be deterministic and bounded on $0 < J \le 1$. Suppose a sufficiently high fixed quantile of F_J is at most $K_0 v(J)$, uniformly over all $J > 0$. Fix $a_0, s, K > 0$. If $I \ge 2$, $A \ge I^{a_0}$, and*

$$v(J) \le KA(J/I)^s, \qquad 1 \le J \le I, \tag{9.3}$$

then there exist $q_0 > 0$ and finite constants C, C' such that

$$\mathbb{P}(F_I > CxA) \le C'x^{-q_0}, \qquad x \ge 1. \tag{9.4}$$

The exponent q_0 can be chosen depending only on d, γ and the fixed geometric scheme; the constants may depend on a_0, s, K, K_0 and the low-intensity bound for v.

Along any family satisfying these hypotheses, the actual graph distances divided by A have a common local upper joining modulus on compact subsets of connected open domains. More precisely, for every $\sigma < s\beta_\eta$, every probability tolerance permits a deterministic constant and errors tending to zero such that nearby points z, w in a fixed compact subset can be joined in its open buffer at normalized cost at most

$$C|z-w|^\sigma + o(1),$$

simultaneously in the endpoints. Consequently the compact internal diameters are tight in these units.

Proof. We give the details separating the quantile input, multiscale connectivity, and terminal errors.

Local open and nice vertices. At level j, put $r_j = 2^{-j}$ and use the deterministic lattice $r_j\mathbb{Z}^d$. A vertex is open when all its incident encoded nearest-neighbor port connections have cost at most

$$B_j = C_1(A2^{-s\beta_\eta j} + 1), \tag{9.5}$$

with their prescribed internal confinement and auxiliary guards. The constant C_1 will be fixed independently of I, j.

Condition on the entire coarse noise above a block's cutoff. If $\overline{H}$ is its average on the fixed enlarged block, the effective scalar intensity is

$$J_{\text{eff}} = Ir^b e^{\gamma\overline{H}}.$$

Off the coarse obstacles, the cutoff extension of $H - \overline{H}$ has a uniformly bounded Cameron–Martin norm for the fine field. Cauchy–Schwarz in the Gaussian density formula therefore bounds the conditional probability of failure of a port test by a fixed constant times the square root of its unshifted failure probability. Conditional Poisson sampling does not change this density. Choose the quantile sufficiently high after fixing the coarse-obstacle threshold. The resulting fine-failure probability is as small as required for Corollary 7.5. Given the coarse noise, tests on disjoint fixed enlargements are independent, by the compact support of the fine noise kernel and the conditional independence of Poisson restrictions.

For two adjacent levels, call a vertex on the coarser lattice *nice* if it is open there and the bounded collection of surrounding fine-lattice vertices is open as well. To estimate its failure probability, use the *common coarser cutoff*; the half-grid ports and both guard patterns were included in F for exactly this purpose. All tests then involve a fixed bounded region of one fine field. A union bound, or the sum in the definition of F, controls their joint failure. Conditional finite-range independence holds between separated nice-vertex tests at this common cutoff. Independence between adjacent scales is neither assumed nor needed. The factor $2^{s\beta_\eta}$ between their cost bounds is absorbed into C_1.

For clarity these probability estimates are first applied with the random cost bounds $K_0 v(J_{\text{eff}})$ in place of (9.5). The global coarse-field event below then makes every successful test open at the deterministic cost bounds (9.5). Thus no conditional independence is claimed after conditioning on a global maximum event.

Coarse maxima and cluster geometry. On the fixed macro buffer, the Gaussian maximum estimate gives constants $\tau > 0, C < \infty$ such that

$$\mathbb{P}\left[\sup H_{r_j} > (\sqrt{2d} + \eta)j\log 2\right] \le C2^{-\tau j}. \tag{9.6}$$

Indeed there are $O(2^{dj})$ points in an r_j-net; the Gaussian tail at level $(\sqrt{2d} + \eta/2)j\log 2$ beats this number by a fixed exponential factor in j. Within each mesh cell the rescaled oscillation has uniformly bounded Gaussian variance and expected supremum, so an additional $(\eta/2)j\log 2$ fails with probability smaller than every fixed exponential in j. On the complementary event,

$J_{\mathrm{eff}} \le I2^{-\beta_\eta j} \le I$. If $J_{\mathrm{eff}} \ge 1$, the past envelope bounds its cost bound by $KA2^{-s\beta_\eta j}$; otherwise the fixed bound on v applies. This proves (9.5), including the corresponding tests using the preceding level's cutoff.

The mixed coarse/fine obstacle estimate bounds the diameter of a closed cluster and of a rough (not nice) cluster by an exponential tail in lattice units. A union bound over $O(2^{dj})$ sites shows that, except with probability $C2^{dj}e^{-cj^2}$, all such clusters have diameter at most a fixed multiple of j^2. These estimates hold simultaneously at all requested levels by a further union bound.

The deterministic lattice argument of [7, Section 4.1] now applies. Declare vertices outside the macro observation box open. For $d \ge 2$, small closed star-clusters give a unique spanning open cluster $\mathcal{O}_j$. Two points of this cluster in a connected region can be joined in a neighborhood of width Cj^2r_j. An infinite nice cluster on the coarser lattice lies in $\mathcal{O}_j \cap \mathcal{O}_{j-1}$, and every point is within Cj^2r_j of that intersection. A connection over that distance uses at most Cj^{2d} lattice edges. This is the part of the argument which replaces planar intersection of crossing traces.

Chaining and the microscopic end. Choose a terminal level $J = \lceil M \log_2 I \rceil$, with M large enough for the microscopic inputs at the desired polynomial failure exponent. Each actual query endpoint is within CJ^2r_J of $\mathcal{O}_J$. The small-ball cell bound attaches its query representative to this nearby deterministic port at cost at most I^u, where $0 < u < a_0/2$ is fixed. Descend through the common vertices of successive clusters. The sum of the edge cost bounds for the two endpoint chains is at most

$$C \sum_{j=j_0+1}^{J} j^{2d} B_j \le CA \sum_{j \ge j_0} j^{2d} 2^{-s\beta_\eta j} + CJ^{2d+1}. \tag{9.7}$$

At a fixed sufficiently large starting level j_0, join the two remaining vertices through $\mathcal{O}_{j_0}$ inside the prescribed trace cube of the target port test. More precisely, use a fixed compact neighborhood of the straight segment between its ports, with positive clearance from that cube's boundary. The open-cluster route and all its local observation buffers lie inside the trace cube when j_* is sufficiently large. A simple lattice path there and all of its selected guard windows stay in the target macro guard region after increasing this fixed j_* once. The path has at most $C2^{dj_0}$ edges. The total undeleted cost is consequently bounded by

$$C(A2^{dj_0} + J^{2d+1} + I^u). \tag{9.8}$$

Only $g \le C(2^{dj_0} + J^{2d+1})$ auxiliary guard occurrences are needed. Initially each selected block trace is constructed in its own guarded diagram. Superimpose only the complete finite guard sets belonging to these selected blocks, together with the fixed macro guard set for the target observable. Inserting their union adds at most $O(g)$ cells to each of $O(g)$ selected traces, hence $O(g^2)$ cells in total. Now all traces use one diagram and concatenate at their exact deterministic ports. Delete the auxiliary guards while retaining the macro guards. The trace-preserving deletion bound is at most $\sum_p L_I(p) + O(g^2)$, and the common screening event bounds this by $gI^u + O(g^2)$. The resulting concatenated trace is in the original target diagram and has the same geometric endpoints. These statements concern only selected guard sets; the union bound establishing their screening was taken in advance over the complete deterministic grids. Repeat for the finite number of target port tests and guard patterns defining F_I.

This conversion uses cell covers throughout. The deterministic grid guards can produce cospherical degeneracies; no assertion about facet adjacency in those intermediate diagrams is needed. Insertion shrinks every old cell and adds just one new cell to a cover. Deletion uses weak empty-ball neighbors, including degenerate witnesses. Convexity of cells and of the trace cube

permits a finite connected cover to be reduced to a cell chain inside the same cube if desired. Only after all guards have been removed do we convert that chain to a graph path in the generic original Poisson configuration. A continuous trace covered by n cells then gives a Delaunay route of at most $n + O(1)$ edges. The vanishing cell mesh keeps its nuclei and edges within a vanishing enlargement of the trace.

A tail uniform in the intensity. Fix $0 < \kappa < 1/(4d)$ and take $j_0 = \lceil \kappa \log_2 x \rceil + j_*$, where j_* is a fixed large geometric constant. First restrict to $1 \le x \le I^K$, with a fixed sufficiently large K, and increase M if necessary so that $j_0 \le J$ outside a bounded range of I, x. Bounded ranges are absorbed into the constants. Because $A \ge I^{a_0}$ and $u < a_0/2$, the bound (9.8) and all deletion errors, divided by A, are at most

$$C\{x^{\kappa d} + x^{2\kappa d} + 1\} \le C' x.$$

Here every expression $(\log I)^m I^{-c}$ has a finite supremum for fixed $m, c > 0$. This explicitly absorbs the low-intensity cost bounds at the deepest levels and the endpoint attachments.

Summing (9.6) for $j \ge j_0$ gives failure probability at most $Cx^{-\kappa\tau}$. The closed- and rough-cluster failures are bounded by $Ce^{-cj_0^2}$, which is smaller than a fixed power of x^{-1} after changing constants. Choose the microscopic failure exponent $B > Kq_0$, where, for example, $0 < q_0 < \min\{\kappa\tau, 1/2\}$. Then $I^{-B} \le x^{-q_0}$ throughout this range of x.

For $x > I^K$, use (9.1) directly:

$$\mathbb{P}(F_I > xA) \le CI^{1-a_0}/x.$$

Choosing K so that $1-(1-a_0)_+/K \ge 1/2$ bounds this by $Cx^{-1/2}$. One can choose K independently of $a_0 > 0$, for example $K = 4$. This proves (9.4) for all x. The choice of q_0 did not use any moment exponent for F.

Local modulus. For two endpoints at Euclidean separation comparable to 2^{-j_0}, the starting-level connection stays in a cube of radius $Cj_0^2 2^{-j_0}$ and needs only Cj_0^{2d} edges, rather than $C2^{dj_0}$. The same descent therefore gives the main normalized cost

$$C \sum_{j \ge j_0} j^{2d} 2^{-s\beta_\eta j} \le C_\sigma 2^{-\sigma j_0}, \qquad \sigma < s\beta_\eta.$$

The remaining terms are bounded uniformly in the endpoints by fixed powers of $\log I$ times I^{u-a_0} or I^{-a_0}, and tend to zero. If the endpoint separation is below the terminal radius, use the uniform small-ball cell bound directly. The probability of a failure at some level $j \ge j_*$ can be made arbitrarily small by taking j_* large; the microscopic failure also tends to zero. The resulting traces stay inside the prescribed open buffer once the starting radius is smaller than its clearance. Remove the macro guards using Lemma 9.1; the original-cell mesh then keeps the corresponding graph paths in that buffer. Finite chains in a connected open domain give tight compact internal diameters. This completes the proof. □

Corollary 9.3 (Fractional-moment records without a target exponent)**.** *Let $c_0 > 0$ be the exponent obtained in Lemma 9.1. Choose $0 < p < \min\{1, q_0/8\}$ and set $b(I) = (\mathbb{E}F_I^p)^{1/p}$. Then b is locally bounded, bounded on $I \le 1$, and for every $0 < a_0 < c_0$,*

$$I^{a_0} \le b(I) \le C(1 + I)$$

for all sufficiently large I. For any $0 < s < a_0$ there is an unbounded sequence of past cone records I_n such that

$$b(J) \le 2b(I_n)(J/I_n)^s, \qquad 1 \le J \le I_n.$$

Along these records, Theorem 9.2 with $A = b(I_n)$ gives uniformly bounded moments of every order $q < q_0$ and uniform integrability of the pth powers. Any joint subsequential limit of $F_{I_n}/b(I_n)$ has a nonzero value with positive probability.

Proof. The upper and local bounds follow from (9.1) and concavity. The proved mesh event (9.2) gives the lower bound. Since $b(I)/I^s \to \infty$, approximate running maxima of this ratio give the stated records; exact attainment of a running maximum is unnecessary. Markov's inequality makes a sufficiently high quantile of F_J at most $K_0 b(J)$, with K_0 depending on the already chosen p. The tail exponent q_0 was chosen before p and is unaffected by this multiplier. Integrating the tail gives the moment bounds. As $\mathbb{E}(F_{I_n}/b(I_n))^p = 1$, uniform integrability preserves this identity in a limit and proves nontriviality. □

Remark 9.4 (What is still not identified). These records give a route to nontrivial subsequential passage laws with local upper compactness in their own units. They do not prove $b(I) = I^{\xi/\gamma+o(1)}$, bounded adjacent-intensity ratios at *all* large intensities, lower comparison with the target metric, or a Weyl rule for a graph limit. Relating these finite-anchor units to the target metric's quantum scaling, completing the enlarged local kernels, and propagating the estimates beyond records remain distinct steps of a full convergence proof.

10 Cox–Palm screening and selected guards

This section makes explicit two interfaces in the discrete estimates. It assumes the local GMC hypotheses of Section 3 and the inverse-mass integral proved in Section 5. It does not assume a metric scaling limit.

Lemma 10.1 (The marked Cox–Palm void bound)**.** *Let p be either a deterministic query or an annealed Palm root, and write $e = 0$ or $e = 1$, respectively. Fix $k \ge 1$ and a shell $r < \max_i |q_i - p| \le 2r$. Suppose that all k inserted points admit empty-ball witnesses against the lower cloud of intensity s. For each insertion tuple, let $\{B_v\}$ be any of the finitely many valid separated subball families from the screening-tree construction. Set $q_i = p + rz_i$, $z_0 = 0$, $\chi_0 = e$, and $\chi_i = 1$ for $i \ge 1$. There are constants depending only on the fixed data and k such that the shell contribution to the kth factorial moment is at most the sum over these finitely many choices of*

$$C_k t^k r^{ka_e - gk^2/2} \int W(z)\, \mathbb{E} \exp\{-c_k s r^{a_e - gk} e^{\gamma\Omega} M_z\}\, dz, \qquad a_e = d + g/2 - eg, \tag{10.1}$$

where $\Omega \sim N(0, \log(1/r) + C)$ is independent of the fixed reference field defining M_z.

Proof. Condition on the field and use the marked Poisson Mecke formula for the k upper-cloud points. Their marks are integrated over $[0, t]$, giving a factor t^k. The GMC multipoint Palm identity supplies the pair factor

$$\prod_{0 \le i < j \le k} e^{g\chi_i\chi_j K(q_i, q_j)}$$

and replaces the background chaos by

$$\prod_{i=0}^{k} e^{g\chi_i K(x, q_i)}\, \mu_h(dx).$$

In the Palm-root case the factor for the root itself is included in the reference Palm law; its deterministic spatial density is bounded on the compact under consideration.

The background lower cloud remains conditionally Poisson with this inserted intensity measure. Inserted upper-cloud points whose marks are at most s impose additional emptiness constraints. Dropping those constraints only increases the probability. Every selected subball avoids all insertion positions by construction. The probability that the background avoids their union is therefore exactly

$$\mathbb{E}\exp\left[-s\int_{\cup_v B_v}\prod_{i=0}^{k}e^{g\chi_i K(x,q_i)}\,\mu_h(dx)\right].$$

It is a void probability for the union, not a product of probabilities for the original witnesses. A finite union bound handles the finite choice of subballs forced by existence of the witnesses.

The change of variables $q_i = p + rz_i$ and the pair density give

$$r^{dk-gk(k-1)/2-egk} = r^{ka_e-gk^2/2}.$$

All bounded covariance-remainder terms can be absorbed into fixed upper and lower multiplicative constants. Since $u \mapsto e^{-u}$ is decreasing, a lower bound on the deterministic insertion weights is used before covariance comparison.

On the rescaled buffer the covariance is $\log(1/r) + \log|x-y|^{-1} + R_r(x,y)$ with uniformly bounded R_r. A fixed reference log field F plus an independent common Gaussian Ω of variance $\log(1/r) + C$ dominates this covariance for sufficiently large fixed C. Kahane comparison for the convex void functional is in the required upper-bound direction. The variance-normalized common Gaussian contributes $r^{g/2}$, spatial volume contributes r^d, and the $k+e$ insertions contribute $r^{-g(k+e)}$. Their product is r^{a_e-gk}, proving (10.1). □

Lemma 10.2 (Summation of the Palm–void bound)**.** *If $\int W(z)\mathbb{E}[M_z^{-k}]\,dz < \infty$, then summing (10.1) over the dyadic shells is bounded by $C_k(t/s)^k$, uniformly over $0 < s \le t$ and the root position in the prescribed compact buffer.*

Proof. Write $T = \log(1/r)$ and $\sigma^2 = T + C$. Gaussian tilting yields

$$\mathbb{E}e^{-c_k s r^{a_e-gk}e^{\gamma\Omega}M} = \mathbb{E}\left[e^{k\gamma\Omega-gk^2\sigma^2/2}e^{-c_k e^{-gkC}sr^{a_e}e^{\gamma\Omega}M}\right].$$

Indeed the positive tilt shifts Ω by $k\gamma\sigma^2$. Multiplication by the shell prefactor cancels $r^{-gk^2/2}$. The resulting integrand is a constant times $(t/s)^k H_k(U_T)$, where

$$H_k(u) = e^{ku}\int W(z)\mathbb{E}e^{-e^u M_z}\,dz, \qquad U_T = \log s - a_e T + \gamma\Omega + c_k'.$$

Tonelli gives

$$\int_{\mathbb{R}} H_k(u)\,du = \Gamma(k)\int W(z)\mathbb{E}[M_z^{-k}]\,dz < \infty.$$

Because $g < 2d$, both possible drifts $a_0 = d + g/2$ and $a_1 = d - g/2$ are positive. Up to a fixed translation, the marginal law of U_T is that of $\sqrt{g}B_u - a_e u$ at $u = T + C$. Its density $f_u(y)$ obeys

$$\int_0^\infty f_u(y)\,du \le a_e^{-1}.$$

For fixed y, $u \mapsto f_u(y)$ is unimodal, and for $u \ge u_0 > 0$ its supremum is at most $(2\pi g u_0)^{-1/2}$. Hence the sum of the densities over the fixed dyadic time lattice is uniformly bounded in y. Integrating against H_k proves the assertion. The argument uses no joint coupling or independence of the shell Gaussians. □

Lemma 10.3 (Trace-preserving guard deletion)**.** *Let P be a locally finite site set and G a finite set of extra sites. Fix a compact trace T covered by finitely many closed cells of the diagram for $P \cup G$. Suppose a fixed buffer or a retained guard net localizes every cell replacement relevant to T. For $p \in G$, let $L_P^0(p)$ count sites of P in that buffer admitting a weak empty-ball witness with p against P. After deleting all of G, the same trace has a cell cover with cardinality at most its original cover cardinality plus*

$$\sum_{p \in G} L_P^0(p) + |G|^2. \tag{10.2}$$

The bound holds for a union of finitely many traces as well, if its initial cover cardinality counts the union of their covering cells.

Proof. Delete the sites one at a time. Outside the former cell of the site p being deleted, all unchanged cells only expand. At a point y of the former cell of p, choose a nucleus q whose new cell contains y. The ball centered at y with q on its boundary is empty of the remaining sites, and it contains p in its closure. If p lies in its interior, shrink this ball by homothety about q until p first reaches its boundary. If p is already on the boundary, retain the ball. The resulting ball is empty of every point of P and has both p and q on its boundary.

Thus every possible replacement nucleus belongs either to the counted set of P-neighbors or to the remaining guards. These cells cover the entire relevant part of the former cell. They are added once to the covering set, irrespective of how often the trace visits that cell. The increase at this deletion is at most $L_P^0(p) + |G|$. Sum over $p \in G$ to obtain (10.2). The trace itself is never changed, and the argument applies without alteration to a union of traces. □

Corollary 10.4 (Guards selected after observing the configuration)**.** *Let $\mathcal{G}_I$ be a deterministic candidate set in a fixed buffer with $|\mathcal{G}_I| \le CI^D$. Assume the uniform rootless screening moments. For every $B, u > 0$, with probability $1 - O(I^{-B})$,*

$$L_{\mathcal{P}_I}^0(p) \le I^u \quad \text{for every } p \in \mathcal{G}_I.$$

On this single event, every configuration-dependent subset $G \subseteq \mathcal{G}_I$ satisfies the deletion bound $|G|I^u + |G|^2$ for every localized trace, simultaneously.

Proof. Choose a moment order $q > (B + D)/u$. Markov's inequality and the uniform rootless moment estimate give

$$\mathbb{P}\left(\max_{p \in \mathcal{G}_I} L_{\mathcal{P}_I}^0(p) > I^u\right) \le C_q I^{D-uq} = O(I^{-B}).$$

The event is established before choosing G. Apply Lemma 10.3 pathwise. No conditional law is assigned to a guard chosen after inspecting the cloud. □

Application to selected anchor networks. Suppose at most g local traces are chosen, and their complete deterministic guarding sets contain at most Cg distinct sites. Insert their union into each chosen diagram. A new site adds at most one cell to a fixed trace cover, so the summed insertion cost is $O(g^2)$. The traces now belong to one diagram and concatenate at their common deterministic ports. Take their union cover, and delete the auxiliary guards with the preceding corollary. The total repair cost is $O(gI^u + g^2)$. Repeated visits to guard cells and choices of traces made after observing the cloud do not increase this bound. Only after deletion, in the generic original Poisson diagram, is the connected cell cover converted to a Delaunay route. The spatial enlargement needed for this final conversion is the vanishing cell mesh, not necessarily the much smaller terminal anchor radius.

11 The collision integral for screening

This supplement concerns only the Euclidean empty-ball estimate in Section 5. No random metric, contact theorem, or metric scaling limit enters. Write $g = \gamma^2 \in (0, 2d)$ and use variance-normalized subcritical GMC for an actual Gaussian field whose covariance, on a fixed enlarged compact buffer, is

$$K(x, y) = \log |x - y|^{-1} + R(x, y), \qquad |R| \le C_R.$$

If a deterministic sampling density is present, the statement below uses the explicit convention $0 < w_- \le w(x) \le w_+ < \infty$ on that buffer. A merely positive bounded measurable density, without a positive lower bound, is not covered by this argument.

Proposition 11.1 (The all-order screening input)**.** *Fix the inner compact set of roots and a sufficiently small fixed screening radius ρ. For the marked Cox coupling $\mathcal{P}_s \subseteq \mathcal{P}_t$, let $L_{s,t}(p)$ count upper sites $q \in B_\rho(p)$ which have a Euclidean open empty-ball witness through p, q against the lower background cloud. Then, for every finite $q \ge 1$,*

$$\sup_p \mathbb{E}[L_{s,t}(p)^q] \le C_q (t/s)^q, \qquad 0 < s \le t.$$

The assertion holds for the original-field query law and for the canonical GMC-rooted law at p. Equivalently, the Palm formulation holds for almost every p with respect to the deterministic mean intensity, and its integrated form gives

$$\mathbb{E} \sum_{p \in \mathcal{P}_t \cap K} L_{s,t}(p)^q \le C_q t (t/s)^q.$$

For each fixed d, g, q the constant is finite. There is no claim of a constant uniform as $g \uparrow 2d$.

The Palm and radial steps are expanded in Section 10. The purpose here is to verify, independently, the finite-choice geometry and the singular collision integral on which those steps depend.

11.1 A deterministic finite family really suffices

Fix an ordered k-tuple of distinct insertion positions $z_0 = 0, z_1, \dots, z_k$, with $c_0 \le \max_i |z_i| \le C_0$, and put $n = k + 1$. All tie choices in the following construction are made lexicographically. Thus the tree, its representatives, scales, and candidate balls are Borel functions of the insertion positions. Let r_v be the largest edge of a minimum spanning tree of the cluster I_v. Split the cluster into the connected components of the threshold-r_v/L graph and iterate. This gives at most k internal vertices, each with at least two children, and

$$\frac{r_v}{L} < |z_i - z_j| \le (n - 1) r_v \quad \text{when } v \text{ is the least common ancestor of } i, j. \tag{S1}$$

The root scale has fixed positive upper and lower bounds.

At an internal node choose an insertion q_v of maximal distance from 0 in I_v. The triangle inequality gives $|q_v| \ge r_v/2$, so every witness ball through $0, q_v$ has radius $R \ge r_v/4$. Let $\delta = 1/100$, $b = \delta/[64(n + 1)]$, and use a fixed Euclidean 1/8-net $\mathcal{U}$ of the unit sphere. If v_0 is the inward unit normal of a witness at q_v, choose $u \in \mathcal{U}$ with $|u - v_0| \le 1/8$. Then

$$\overline{B}(q_v + \delta r_v u, \delta r_v/4) \subset \text{int(witness)}. \tag{S2}$$

Indeed, writing $a = \delta r_v$, we have $a/R \le 4\delta$ and $u \cdot v_0 \ge 127/128$. The inequality $|au - Rv_0| + a/4 < R$ follows after squaring from $u \cdot v_0 > 1/4 + 15a/(32R)$.

Inside the ball in (S2), place $n+1$ closed balls of radius $2br_v$, whose centers are $8br_v$ apart on a fixed coordinate line centered at its center. Their total span fits strictly in the ball in (S2). The closed balls are pairwise disjoint; at most n of them contain one of the n insertion positions. For an unoccupied one retain its concentric radius-br_v ball B_v. Consequently,

$$\operatorname{dist}(B_v, \{z_0, \dots, z_k\}) \ge br_v, \qquad B_v \subseteq B(q_v, 2\delta r_v). \tag{S3}$$

Per node there are at most $|\mathcal{U}|(n+1)$ candidates. The total number of choices is therefore at most $[|\mathcal{U}|(n+1)]^k$, uniformly over positions and field realizations. Existence of all witnesses implies emptiness of at least one valid family. We sum over all valid families after fixing the positions; we do not condition on a field-selected direction or family.

Choose $L \ge 4\delta/b$. If w is a proper descendant of v, then $r_w \le r_v/L$ and q_w is an insertion. Thus

$$\operatorname{dist}(B_v, B_w) \ge br_v - 2\delta r_w \ge br_v/2.$$

For incomparable nodes v, w with least common ancestor u, their representatives are in different threshold-r_u/L components and $r_v, r_w \le r_u/L$. Hence

$$\operatorname{dist}(B_v, B_w) > r_u/L - 2\delta(r_v + r_w) \ge r_u/(2L). \tag{S4}$$

This is the separation needed for the Gaussian comparison. No planar topology or disjointness of the original witnesses is used.

11.2 Tree charges and a valid Gaussian comparison

Let $e = 0$ for a query and $e = 1$ for a GMC root; set $\chi_0 = e$ and $\chi_i = 1$ for $i \ge 1$. On a fixed normalized buffer choose a genuine reference log field F with bounded covariance remainder and its subcritical GMC. For any one of the valid families define

$$W(z) = \prod_{i<j} |z_i - z_j|^{-g\chi_i\chi_j}, \qquad M_z = \int_{\cup_v B_v} \prod_i |x - z_i|^{-g\chi_i}\, \mu_F(dx).$$

Every B_v is separated from all insertions, so this last measure is finite and strictly positive for every fixed distinct tuple. The needed claim is

$$\int W(z)\mathbb{E}[M_z^{-k}]\, dz_1 \cdots dz_k < \infty. \tag{S5}$$

Put $c = d/g + 1/2 > 1$ and $c_v = \sum_{i \in I_v} \chi_i$. At proper internal nodes let $P_v = (c_v - c)_+$, at the root put $P_{\text{root}} = k$, and at leaves put $P_v = 0$. Then

$$s_v = P_v - \sum_{w \text{ child of } v} P_w \ge 0, \qquad \sum_v s_v = k. \tag{S6}$$

For a proper node, this follows from $\sum_i (a_i - c)_+ \le (\sum_i a_i - c)_+$ for nonnegative a_i. At the root the child sum is at most $(k + e - c)_+ \le k$. Leaves have charge at most one, below c.

Let $\theta_u = r_u/r_{\text{par}(u)}$ for a proper internal node. For $x \in B_v$, successive ancestor diameter bounds imply

$$\prod_i |x - z_i|^{-g\chi_i} \ge C^{-1} A_v, \qquad A_v = \prod_{u \preceq v,\ u \ne \text{root}} \theta_u^{-gc_u}. \tag{S7}$$

The inequality uses upper bounds on distances; insertion-free separation is not used in the wrong direction here. Disjointness and weighted AM–GM, omitting factors with $s_v = 0$, now give

$$M_z^{-k} \le C \prod_v (A_v \mu_F(B_v))^{-s_v}. \tag{S8}$$

Choose C_* larger than the fixed root-scale bound. Define jointly Gaussian tree variables by a root variance $\log(C_*/r_{\mathrm{root}})$ and independent edge variances $\log(r_{\mathrm{par}(v)}/r_v)$. Then $\mathrm{Cov}(G_v, G_w) = \log(C_*/r_{\mathrm{lca}(v,w)})$. On each B_v add an independent copy of a reference log field rescaled by r_v, independent also of the tree variables. Equation (S4) and bounded remainders show that a fixed sufficiently large C_* makes this Gaussian field dominate the covariance of F on the disjoint union of balls. Within a single ball the two logarithmic scale terms cancel exactly. Across balls (S4) bounds the original covariance by the corresponding tree covariance. The comparison field is positive semidefinite by its construction.

The relevant form of Kahane comparison can be derived using only bounded Laplace functionals. For deterministic $t_v > 0$, apply the convex scalar function $x \mapsto e^{-x}$ to the base measure $\sum_v t_v \mathbf{1}_{B_v}(x)\,dx$. Regularize and then pass to the limit by bounded convergence. Integrate the resulting inequality against $\prod_v t_v^{s_v - 1}/\Gamma(s_v)$ for the indices $s_v > 0$. Tonelli gives the product negative-moment comparison. No assumed uniform negative moment of a regularized mass is required for this passage.

For the comparison field the masses have the exact form, up to fixed deterministic ball-shape factors,

$$\mu_Y(B_v) = C_*^{-g/2} r_v^{d+g/2} e^{\gamma G_v} Z_v, \tag{S9}$$

where the independent Z_v are masses of fixed reference balls and have all negative moments. By (S6), the sum of the exponents s_w over all descendants of v, including v, is P_v. Taking Gaussian expectations in (S8), every proper tree edge thus contributes exactly

$$\theta_v^{P_v(gc_v - d - g/2) - gP_v^2/2} = \theta_v^{\frac{g}{2}(c_v - c)_+^2}. \tag{S10}$$

The root Gaussian contributes only a bounded constant, since its scale is bounded above and below.

11.3 Collision integrability, including high orders

For a proper cluster of $m \ge 2$ charged leaves, the configuration volume, pair-density singularity, and (S10) give the exponent

$$E_m = d(m-1) - \frac{g}{2} m(m-1) + \frac{g}{2}(m-c)_+^2. \tag{S11}$$

If $m \le c$, then $E_m = (m-1)(d - gm/2) > 0$, since $gm \le gc = d + g/2 < 2d$. If $m > c$, direct expansion cancels both powers of m and yields

$$E_m = \frac{(2d-g)^2}{8g} > 0. \tag{S12}$$

For a cluster containing the uncharged root and $j \ge 1$ charged leaves, the exponent is instead

$$E_j^0 = dj - \frac{g}{2} j(j-1) + \frac{g}{2}(j-c)_+^2 > 0. \tag{S13}$$

Below threshold the assertion is immediate from $j \le c$; above threshold its value is $gc^2/2$.

Here is the complete volume bookkeeping. Give every cluster a representative inherited from one child, with the root representative equal to $z_0 = 0$. At a node with ℓ_v children introduce the $\ell_v - 1$ differences of the remaining child representatives from the chosen one. Over the whole tree these are k vectors. The linear map from them to $(z_1, \dots, z_k)$ has determinant ± 1, coordinate by coordinate. Each difference at v has norm at most $(n-1)r_v$. On a dyadic scale assignment this bounds the configuration volume by

$$C \prod_v r_v^{d(\ell_v - 1)}.$$

The sum of $\ell_w - 1$ over the internal descendants of v is $|I_v| - 1$. Rewriting in edge ratios therefore gives exactly $d(|I_v| - 1)$ at a proper node. Equation (S1) gives $-g\binom{c_v}{2}$ for the pair factor at that node. Combining these with (S10) gives (S11) or (S13), with no missing translation coordinate or multifurcation factor.

There are finitely many tree shapes, labels, representative choices, and ball-choice indices. Quantize each proper edge ratio dyadically. The preceding positive exponents bound its scale sum by a geometric series. Independent summation over all edge assignments only enlarges the permitted nested-scale set. This proves (S5). In particular, arbitrarily large fixed k does not require $\mathbb{E}[\mu_F(U)^k] < \infty$. It requires only the fixed reference-ball negative moments and the strict margin (S12).

11.4 High-order Palm identities and the direction of deletion

One must not assume that the kth ordinary moment of the total GMC mass is finite to apply this argument. Apply the GMC multipoint identity first on sets of insertion tuples with a fixed positive mutual separation. There the off-diagonal pair density is bounded. The GMC identity with bounded nonnegative test functionals is valid on such sets, and the marked Mecke formula applies without an integrability hypothesis. Increase these sets to all distinct tuples. Atomlessness removes diagonals and nonnegative monotone convergence gives the identity, potentially with both sides initially infinite. Estimate (S5) and the radial bound establish finiteness afterwards.

After inserting the upper-cloud sites and, if appropriate, the root, the background intensity is

$$s \prod_i e^{g\chi_i K(x,q_i)}\, \mu_h(dx).$$

An inserted upper site with mark at most s could invalidate one of the original witnesses. Discarding this constraint enlarges the event; it is the correct direction for an upper bound. Moreover all selected subballs avoid every insertion by (S3), so none of these deterministic inserted sites affects their void probability. The upper bound is the single union-void functional

$$\mathbb{E} \exp\left[-s \int_{\cup_v B_v} \prod_i e^{g\chi_i K(x,q_i)} \mu_h(dx) \right],$$

followed by the finite choice union bound. It is neither a product of witness probabilities nor a conditional independence claim.

The shell comparison and Gaussian tilt of Section 10 now apply. Their two possible radial drifts are $a_0 = d + g/2$ and $a_1 = d - g/2$, both strictly positive. The sum over dyadic shells is bounded by the potential density of a Brownian motion with this negative drift, whose continuous potential is at most $1/a_e$. Thus every factorial moment is bounded by $C_k(t/s)^k$. Ordinary moments follow by the factorial expansion, and real orders by Jensen. This proves Proposition 11.1.

What the verification does and does not supply. It supplies uniform annealed query and rooted moments at fixed subcritical g, including the deterministic candidate-grid union bound used for adaptively selected guards. It does not supply uniform conditional-on-the-field moments, all-order moments of the total point count in a fixed domain, a bound uniform over every continuum-chosen query point, or uniformity at criticality. None of those stronger statements is used in the selected-guard lemma.

12 Finite-window stabilization, ties, and discrete local readouts

This supplement isolates the deterministic geometry needed before a positive limiting metric is identified. It concerns confined walks and cell covers. It does not assume that an unrestricted shortest graph path is locally determined.

For a compact set K, write $K_r = \{x : \mathrm{dist}_{\mathrm{Euc}}(x, K) \le r\}$.

12.1 The whole-space cloud is well defined

Lemma 12.1 (Global Poisson geometry at every positive intensity)**.** *For canonical subcritical whole-space logarithmic GMC with a fixed smooth unit-mass pin, almost surely every nonempty open cone has infinite mass. In the common marked Poisson construction, almost surely, simultaneously for every real $I > 0$, the cloud is nonempty, locally finite, and in general position; all its Euclidean Voronoi cells are bounded and its facet-adjacency graph is connected. Every compact set meets only finitely many cells.*

Proof. Choose a fixed Euclidean ball B in an annulus away from the origin whose positive dilates lie in a prescribed cone. Set $r_n = 2^n$, $X_n = h_{r_n}(0)$ with the scaled smooth pin, and $g_n = h(r_n\,\cdot) - X_n$. The zero-mass covariance identity gives $g_n \overset{\mathrm{law}}{=} h$ and $\mathrm{Var}\, X_n = \log r_n + O(1)$. Canonical GMC scaling gives

$$\mu_h(r_n B) = r_n^{\gamma Q} e^{\gamma X_n} \mu_{g_n}(B), \qquad \gamma Q = d + \gamma^2/2 > 0.$$

The standard negative moments of subcritical GMC on a fixed ball [9, 10] and a Gaussian tail bound show that, for any fixed $0 < \varepsilon < \gamma Q$,

$$\sum_n \mathbb{P}\Big[X_n < -\frac{\varepsilon}{2\gamma} \log r_n \text{ or } \mu_{g_n}(B) < r_n^{-\varepsilon/2}\Big] < \infty.$$

Independence of these variables is unnecessary. Eventually $\mu_h(r_n B) \ge r_n^{\gamma Q - \varepsilon} \to \infty$. Thus the cone has infinite mass. Intersect over a countable angular basis; every nonempty open cone contains a basis cone. Translated cones have the same conclusion by taking a narrower tail cone based at the origin.

Conditional on the field, every such cone contains infinitely many Poisson points at each positive rational intensity. Monotonicity of the marked cloud extends this simultaneously to every $I > 0$. Local finiteness follows from the Radon property, first at integer intensities and then at all smaller intensities.

For general position, restrict to a compact set and to tuples with pairwise separation at least $1/m$. Their annealed factorial measure has a bounded density by the off-diagonal GMC insertion formula. The affine-degeneracy locus for $d+1$ points and the cospherical locus for $d+2$ points are Lebesgue null. A countable union over compact sets, m, and integer intensity bounds excludes these configurations simultaneously in every subcloud. The nonatomic intensity also makes the cloud simple. This argument uses no unrestricted high GMC moment.

Fix any nucleus p and any $I > 0$. Infinitely many nuclei in every angular cone allow finitely many choices $q_1, \dots, q_N$ such that the vectors $q_i - p$ positively span $\mathbb{R}^d$. For example, choose sufficiently distant points in small cones around the positive and negative coordinate directions. The cell of p is contained in the intersection of the half-spaces

$$2(q_i - p) \cdot (x - p) \le |q_i - p|^2.$$

This finite intersection is bounded: a nonzero recession vector would have nonpositive scalar product with every positively spanning normal. Thus every cell is bounded.

Fix one site q. If a cell with nucleus p meets a compact K at x, then $|p - x| \le |q - x|$. The right side is bounded on K, so all such p lie in a bounded set, containing only finitely many sites. This proves local finiteness of the diagram. A compact polygonal curve between interiors of two cells can now be perturbed to avoid their locally finite faces of codimension at least two. It crosses only finitely many facets and supplies a finite facet-adjacent cell chain. Hence the graph is connected. In dimension one the same assertion is the ordering of consecutive sites. These last statements are deterministic consequences of the simultaneous cone and general-position events, so require no uncountable intersection over sites or intensities. □

12.2 Local filling and stabilization

Lemma 12.2 (An explicit filling-to-stabilization implication)**.** *Let P be locally finite, let W be a bounded open observation window, and suppose $K_{3\rho} \subset W$. Fix $0 < \delta < \rho/32$. Assume that every open Euclidean ball of radius δ with center in $K_{2\rho}$ contains a point of $P \cap W$.*

Every open ball empty of $P \cap W$ whose closure meets K_ρ has radius less than 4δ. Consequently:

1. *every closed Voronoi cell meeting K lies in $K_{5\delta}$ and has diameter at most 8δ;*
2. *every Delaunay edge whose straight segment meets K has length at most 8δ, even if neither endpoint belongs to K;*
3. *the entire cells meeting K, their incident adjacencies, and all Delaunay edges whose segments meet K agree for P and $P \cap W$.*

The assertions apply to both diagrams, and they are uniform over all walks and all query points inspected in K.

Proof. Let $B(c, R)$ be empty and take $x \in \overline{B}(c, R) \cap K_\rho$. If $R \ge 4\delta$ and $|c - x| \ge 2\delta$, put $y = x + 2\delta(c - x)/|c - x|$. Then $|y - c| \le R - 2\delta$, so $B(y, \delta) \subset B(c, R)$. If $|c - x| < 2\delta$, use $y = c$ instead. In both cases $y \in K_{2\rho}$, contradicting the filling assumption.

If $x \in K$ belongs to the cell of p, a point in $B(x, \delta)$ gives $|p - x| < \delta$. For every y in that cell, the ball $B(y, |y - p|)$ is empty and its closure contains $p \in K_\rho$. Thus $|y - p| < 4\delta$. This gives the cell assertions. A Delaunay edge has an empty witness ball with both endpoints on its boundary. Convexity puts the whole edge segment in the closure of this ball. If the segment meets K, the empty-ball assertion bounds its radius by 4δ and its diameter by 8δ.

At every point of K_ρ there is a point of $P \cap W$ at distance less than δ, whereas every point outside W is at distance greater than ρ. Hence the two diagrams have exactly the same nearest-site sets throughout K_ρ. Every cell meeting K is already contained in this region, for either diagram, so its entire cell and every adjacency incident to it coincide.

For an edge whose segment meets K, its endpoints lie in $K_{8\delta}$. The empty-ball argument applied at such an endpoint bounds its entire cell within $K_{12\delta} \subset K_\rho$. The same nearest-site comparison therefore identifies both endpoint cells and their adjacency. This works in both directions, ruling out new edges as well as lost ones. □

Why a long exterior chord is not a loophole. An edge with both nuclei outside an inspected compact can still cross it. Controlling only cells with nuclei inside the compact would not rule out this possibility. The empty witness ball for the entire segment is the additional argument in the preceding lemma. It also bounds the first edge by which a graph path exits a fixed neighborhood, without any compact-containment assumption on the remainder of that path.

Corollary 12.3 (Probabilistic stabilization and polynomial units)**.** *For a nonatomic full-support Radon measure and the marked Poisson coupling, the filling assumption of Lemma 12.2 holds eventually almost surely for every fixed K, ρ, δ. Under the stated local subcritical GMC moment assumptions it holds, with probability tending to one, for $\delta = I^{-a}$ and some deterministic $a > 0$. Thus a cell cover joining one fixed separated pair of ports needs at least $I^{a'}$ cells with probability tending to one, for any fixed sufficiently small $a' > 0$. The same lower bound holds after insertion of any fixed deterministic guard net.*

Proof. A finite $\delta/4$-net of $K_{2\rho}$ supplies finitely many positive-mass balls whose filling implies the required event. Their Poisson arrival times are finite almost surely. The countable intersection over a deterministic sequence of radii is legitimate.

For the quantitative assertion, use

$$\mathbb{E}\mu(B_r(z))^{-1} \le Cr^{-(d+g)}, \qquad g = \gamma^2.$$

A grid of $O(r^{-d})$ balls has filling failure at most $CI^{-1}r^{-(2d+g)}$, using $e^{-x} \le Cx^{-1}$. Take r a fixed multiple of I^{-a} with $0 < a < 1/(2d+g)$. The deterministic lemma bounds every relevant cell diameter by CI^{-a}. Inserting guards does not enlarge cells of original sites, and every guard cell also has empty nearest-site balls against the original Poisson cloud. It has the same bound.

If N closed cells cover a continuous trace between two separated ports, the cells that meet the trace contain a connected finite intersection chain between the endpoint cells. A simple chain uses each cell at most once. Its diameter is at most the sum of the cell diameters, so $N \ge cI^a$. Any $a' < a$ gives the stated event for large I. □

Lemma 12.4 (General position for singular sampling measures)**.** *Suppose μ is nonatomic and its annealed multipoint measures, restricted to distinct points separated by a fixed positive distance in a compact set, have locally bounded densities with respect to Lebesgue measure. Almost surely every Euclidean sphere and every affine hyperplane has zero μ-mass. A Poisson cloud sampled from $I\mu$ is in Euclidean general position almost surely, simultaneously for every intensity in the marked coupling.*

Proof. The locus of $d+2$ cospherical points is Lebesgue null. On each separated compact configuration set, its expected μ^{d+2}-mass is therefore zero. A countable exhaustion gives zero product mass on every separated cospherical tuple almost surely. If some sphere had positive nonatomic μ-mass, its $(d+2)$-fold product would give positive mass to distinct tuples, and therefore to tuples separated by $1/n$ in some compact set, for some n. This is a contradiction. The same argument on affine-dependence loci gives the hyperplane assertion and the lower-dimensional degeneracy exclusions. All these are simultaneous statements about the random measure; one does not take an uncountable intersection over spheres.

Conditional Poisson factorial moments then show that every finite cloud has no degenerate tuples. Intersect over integer intensities and bounded windows. Every tuple present at a finite real intensity is contained in one of these clouds. For GMC the required separated density property is precisely the multipoint Palm density; possible divergence near the diagonal is irrelevant here. □

General position does not remove arbitrary query ties. Even a generic cloud has Voronoi facets. A query selected after observing the cloud can lie on one. The correct conclusion is that the incident nuclei at a query form a Delaunay simplex: there are at most $d+1$ of them, they are affinely independent, and every two are joined by an edge. One way to see the last statement is to perturb the query within the bisector of the two chosen nuclei so that all the other incident inequalities become strict; affine independence permits this perturbation. Local finiteness gives a strict margin from every other nearby nucleus.

Thus two choices of representative nucleus differ by at most one hop. Pair distances differ by at most two hops. If the query lies in a fixed compact subset of an open confinement domain, the mesh bound puts this correcting edge inside that domain eventually. No such claim is made exactly on its boundary without an open margin. In units $B_I \to \infty$, the ambiguity is at most $2/B_I$ and vanishes uniformly over all queries with that margin. Deterministic guard diagrams may be degenerate; the argument does not use their edge graphs. It retains closed-cell covers there, and converts connected covers to graph walks only after removing guards in the generic original cloud.

Lemma 12.5 (Measurable finite-window record maps)**.** *For a fixed finite observation window, guarded finite port costs, confined finite graph walks, and their closed collections of ordered marked records are Borel functions of the point configuration. On the filling event of Lemma 12.2, the records inspected in K agree exactly with those of the original infinite diagram. Consequently each fixed finite projection of the asymptotic record array has a local finite-window version with error probability tending to zero.*

Proof. Enumerate the finite configuration measurably. Its Voronoi cells are given by the inequalities

$$2x \cdot (q-p) \le |q|^2 - |p|^2.$$

Nonempty intersections with a fixed polyhedral chart or open trace cube are finite real semialgebraic feasibility conditions in the site coordinates; their parameter sets are Borel. This proves measurability of adjacency, even if one instead uses the semialgebraic definition of a common facet.

For a guarded port observable, each cell intersected with its open trace cube is convex. A fixed selected family covers a connection between the ports exactly when its intersection graph has a chain from a port-containing cell to a cell containing the other port. This equivalence follows in one direction from connectedness of the trace and in the other by concatenating line segments in the convex pieces. Taking the minimum number of selected cells over the finitely many subsets is a Borel operation. It requires no geodesic selector and remains valid in a degenerate guard diagram.

Finite walks in an enumerated configuration form a countable set. Their marked vertices and compact polygonal subranges are Borel; confinement in a compact polyhedral chart is also Borel. Use rational upper cost bounds. The closure of these records hits an open set of the compact record space exactly when one of the countably many records hits it. This gives the claimed measurable random closed collection. The deterministic stabilization lemma identifies every adjacency and query cell used by a confined walk, uniformly over the countable set of such walks.

For finitely many retained charts and finitely many fixed positive intensity multipliers, use the union of their observation windows and their filling events. Their actual intensities all tend to infinity. This proves the finite-projection assertion. It does not require one finite observation event that works at once for all countably many multipliers, including arbitrarily small ones. □

The distinction between physical and normalized scales. The mesh statement uses the actual point intensity, while the one-hop error uses the chosen denominator. In a moving-scale

diagonal both must be checked. If the base denominator obeys $B_n \geq I_n^{a_0}$ and a deterministic label $v_j > 0$ is first selected and then $n = n(j)$ is taken sufficiently large, it is permissible to require

$$v_j I_{n(j)} \longrightarrow \infty, \qquad |\log v_j| = o(\log I_{n(j)}), \qquad v_j^{\alpha} B_{n(j)} \geq (v_j I_{n(j)})^{a_0/3}.$$

The last inequality follows for large $n(j)$ from the first two and the fixed positive exponents. Thus a single edge and every fixed-power subpolynomial terminal repair vanish in the moving units. This is a diagonal choice after selecting v_j; it is not an assertion about an arbitrary preassigned sequence of vanishing multipliers.

Scope and a density convention. These lemmas establish local stabilization and measurability before noncollapse. They do not identify unrestricted global distances with internal distances. That later identification still requires the positive local comparison and the reference metric's escape bound. Likewise, the rough exponent in the corollary is only a polynomial lower bound, not the calibrated exponent ξ/γ. Whenever a deterministic sampling density is said to change only constants, it is required to be bounded above and bounded away from zero on the buffers in use. The canonical-to-variance-normalized densities in this argument have these bounds on fixed compact buffers.

13 A mass–diameter estimate and the metric size of empty cells

For the reference fixed in Section 2, use [8, Propositions 4.7 and 4.10]: normalized crossing costs and the costs of joining arbitrary shell-crossing paths have moments of every positive order. Together with negative GMC moments, these give an all-moment mass–diameter bound for Euclidean balls. The Poisson argument then converts this bound into a quantitative metric size for the actual Voronoi cells. Set

$$\alpha = \xi/\gamma, \qquad \Delta = 1/\alpha, \qquad s_r(z) = r^{\xi Q} e^{\xi h_r(z)}.$$

The field and chaos use the common canonical convention:

$$\mu_h(r\,du + z) = s_r(z)^{\Delta} \mu_{h(r\cdot+z)-h_r(z)}(du)\,.$$

The field $h(r\cdot+z) - h_r(z)$ has the pinned unit-scale law. These are marginal identities for the field and measure, not assertions about the law of a rescaled metric. Independence of the coarse average is not used. The metric input is exactly the uniform moment estimate (13.1) with this deterministic scale normalization.

13.1 A deterministic shell-chain construction

Write $A_{a,b}(z) = B_b(z)\setminus\overline{B}_a(z)$, with Euclidean balls. Let $C_r(z)$ be the minimum reference length of a path crossing from $\partial B_r(z)$ to $\partial B_{4r}(z)$, inspected in $A_{r/2,8r}(z)$. Clip a near minimizer to its first such crossing; its trace then lies in $\overline{A}_{r,4r}(z)$. Let $A_r^{(1)}(z)$ and $A_r^{(2)}(z)$ be the shell-connection quantities from [8, Proposition 4.10], for $A_{5r/4,7r/4}(z)$ and $A_{r,2r}(z)$. Thus each is a supremum, over two crossing paths, of their internal set-to-set distance in the indicated shell. Put $T_r = C_r + A_r^{(1)} + A_r^{(2)}$. The reference assumptions imply

$$\sup_{r,z} \mathbb{E}[(T_r(z)/s_r(z))^q] < \infty, \qquad q > 0. \tag{13.1}$$

Using slightly larger open inspection shells handles their boundary endpoints, without changing any of the fixed shapes.

Fix $\delta > 0$ small in terms of d. For $r_j = 2^{-j} r_0$ choose $z_j(p) \in \delta r_j \mathbb{Z}^d$ with $|z_j(p) - p| \le r_j/100$. Then $|z_j - z_{j+1}| \le 3r_j/200$. A crossing of $A_{r_j,4r_j}(z_j)$ and a crossing of $A_{r_{j+1},4r_{j+1}}(z_{j+1})$ both cross $A_{5r_j/4,7r_j/4}(z_j)$. They can therefore be joined at cost at most $2A^{(1)}_{r_j}(z_j)$, with an arbitrarily small allowance. If $\sum_j T_{r_j}(z_j) < \infty$, near minimizers and these connecting paths give a path from p to the outer endpoint of the first crossing, of length at most

$$C \sum_{j \ge 0} T_{r_j}(z_j(p)). \tag{13.2}$$

Indeed the inner endpoints tend to p in Euclidean distance, the remaining length tends to zero, and the reference metric induces the Euclidean topology. This also justifies the infinite concatenation.

To connect $p, q \in B_R(c)$, take a dyadic $r_0 \in [16R, 32R]$ and a grid center z_B within $r_0/100$ of c. Both spokes cross $A_{r_0/4,r_0/2}(z_B)$; its shell-connection cost joins them. All the paths stay in a fixed enlargement $B_{CR}(c)$. The construction uses shell connections, not intersections of unrelated paths, and works in every $d \ge 2$.

13.2 Uniform inner-ball mass controls diameter

Theorem 13.1 (Mass–diameter bound at small mass)**.** *Fix a compact buffer K, $\theta \in (0,1)$, and a sufficiently small fixed $r_* > 0$. For every $\varepsilon \in (0, \alpha)$ there is a nonnegative random variable H_ε, with*

$$\mathbb{E} H_\varepsilon^q < \infty \qquad \text{for every } q > 0,$$

such that simultaneously for all $z \in K$ and $0 < r \le r_*$ with $\mu_h(B_r(z)) \le 1$,

$$\operatorname{diam}_{D_h}(B_{\theta r}(z)) \le H_\varepsilon \, \mu_h(B_r(z))^{\alpha - \varepsilon}. \tag{13.3}$$

The same bound holds for the internal diameter in one fixed enlargement $B_{Cr}(z)$.

Proof. We work throughout in a fixed enlargement of K, containing all the inspection shells. Fix $\eta > 0$. At dyadic scales $\rho = 2^{-j}$ use the grids $\delta\rho\mathbb{Z}^d$. The following estimates hold with one random $A_\eta \ge 1$ having moments of every positive order:

$$T_\rho(x) \le A_\eta \rho^{-\eta} s_\rho(x), \tag{13.4}$$

$$\mu_h(B_\rho(x)) \ge A_\eta^{-1} \rho^\eta s_\rho(x)^\Delta, \tag{13.5}$$

$$e^{\xi |h_\rho(x) - h_{\rho'}(x')|} \le A_\eta \rho^{-\eta}. \tag{13.6}$$

The last assertion is required only when ρ/ρ' belongs to a fixed compact interval and $|x - x'| \le L\rho$, for a fixed L; both points belong to their respective grids. There are $O(\rho^{-d})$ tests at each scale, including these pairs. Equation (13.1), all negative moments of the normalized GMC ball mass, and Gaussian exponential moments prove these statements by a moment union bound. For the Gaussian comparison, the difference has uniformly bounded variance: its test measure has total mass zero, comparable averaging radii, and separation $O(\rho)$. At any desired sufficiently large moment order q, the union-bound series is bounded by $C_q \sum_j 2^{-(\eta q - d)j}$. Lower moment orders follow from higher ones.

Choose $0 < \kappa_0 < \xi(Q - \sqrt{2d})$. There is also a random $B_0 \ge 1$ having some positive moment $q_* > 0$, such that on the same grids

$$s_\rho(x) \le B_0 \rho^{\kappa_0}. \tag{13.7}$$

To see this directly, uniformly on a compact buffer, $\operatorname{Var} h_\rho(x) = \log(1/\rho) + O(1)$. Hence the moment union bound converges whenever

$$q(\xi Q - \kappa_0) - \frac{\xi^2 q^2}{2} > d.$$

Such a q exists by the choice of κ_0. In what follows take $\kappa = \kappa_0/2$ and $\eta < \kappa_0/2$.

Fix $B = B_r(z)$, put $m = \mu_h(B) \le 1$, and let $p \in B_{\theta r}(z)$. For each grid center used in the spoke from p, at scale $\rho \le 32r$, one can find a grid ball $B_{\rho'}(x') \subset B$ with $\rho' \asymp_\theta \rho$ and $|x - x'| = O_\theta(\rho)$. For ρ below a fixed multiple of $(1-\theta)r$, take a small ball centered near p; for the remaining bounded range of scale ratios use a fixed interior ball of B. The same construction applies to the macroscopic joining shell. Equations (13.5)–(13.6) then imply, after increasing deterministic constants,

$$s_\rho(x) \le C A_\eta^{1+\alpha} m^\alpha \rho^{-(1+\alpha)\eta}.$$

Combining with (13.4) and (13.7) gives, with $e = (2+\alpha)\eta$ and an all-moment random $A \ge 1$,

$$T_\rho(x) \le C A \min\{m^\alpha \rho^{-e},\ B_0 \rho^\kappa\}. \tag{13.8}$$

The elementary dyadic summation, split at $\rho = (m^\alpha/B_0)^{1/(\kappa+e)}$, is

$$\sum_{\rho \in 2^{\mathbb{Z}}} \min\{m^\alpha \rho^{-e}, B_0 \rho^\kappa\} \le C_{e,\kappa} B_0^{e/(\kappa+e)} m^{\alpha\kappa/(\kappa+e)}. \tag{13.9}$$

Apply the shell-chain construction and its one macroscopic join. This bounds the diameter by the right side of (13.9), times CA.

Here is the reason that the claimed constant has *all* moments, despite the limited positive moment range for raw metric diameters. Given any desired moment order $q > 0$, choose η so small that

$$\frac{\alpha\kappa}{\kappa+e} \ge \alpha - \varepsilon, \qquad \frac{2qe}{\kappa+e} < q_*.$$

Since $m \le 1$, the mass power in (13.9) is at most $m^{\alpha-\varepsilon}$. The constant $CAB_0^{e/(\kappa+e)}$ has a finite qth moment by Cauchy–Schwarz. Taking H_ε to be the least uniform constant for rational balls with an inner ratio $\theta' \in (\theta, 1)$, this argument at each q proves all its positive moments. That supremum is measurable, using a countable dense set for each inner diameter. Given an arbitrary target ball $B_r(z)$, approximate it from inside by rational balls $B_{r_n}(z_n) \subset B_r(z)$ with $(z_n, r_n) \to (z, r)$. For large n, $B_{\theta r}(z) \subset B_{\theta' r_n}(z_n)$, and the approximating ball has no larger mass. This transfers the bound with the same random constant; a fixed extra enlargement handles the internal-diameter version. No simultaneous assertion about masses of random spheres is needed. □

13.3 Sharp metric mesh for the Euclidean Poisson cells

Proposition 13.2 (Empty-cell metric diameter)**.** *Let $\mathcal{P}_I$ have conditional intensity $I\mu_h$ and let $\mathcal{V}_I$ be its Euclidean Voronoi tessellation. For every compact K, every $\varepsilon > 0$, and every $B > 0$,*

$$\mathbb{P}\left[\max_{C \in \mathcal{V}_I : C \cap K \ne \varnothing} \operatorname{diam}_{D_h} C > I^{-\alpha+\varepsilon}\right] \le C_{K,\varepsilon,B} I^{-B} \tag{13.10}$$

for all sufficiently large I.

Proof. Fixed-radius filling, estimated using all negative GMC ball moments, confines every cell meeting K to one fixed larger buffer and gives it a sufficiently small Euclidean diameter, with failure $O(I^{-B})$ for every prescribed B. The conditional empty-ball VC bound gives, with the same type of failure estimate,

$$\mu_h(B) \le \delta_I := C_B \log I/I \quad \text{for every relevant empty Euclidean ball } B. \tag{13.11}$$

To make this statement quantitative without requiring a high moment of the total mass, restrict first to $\mu_h(V) \le I^{B+1}$, whose failure is $O(I^{-B-1})$ by the first GMC moment. The VC polynomial prefactor is then absorbed by increasing C_B.

We also use the elementary uniform GMC upper power: for some $a > 0$ there is $K_a \ge 1$ with some finite positive moment such that

$$\mu_h(B_r(z)) \le K_a r^a \tag{13.12}$$

on the fixed buffer for every $0 < r \le r_*$. One may take any $a < (\sqrt{d} - \gamma/\sqrt{2})^2$; moment estimates and a dyadic net give a constant with a positive moment. Consequently $\log K_a$ has moments of every positive order.

Let C have nucleus p and take $x \in C$. Set $x_j = p + 2^{-j}(x - p)$ and $R_j = |x_j - p|$. Convexity puts every x_j in C, and $B_{R_j}(x_j)$ is empty of Poisson sites. Both x_j and x_{j+1} belong to $B_{3R_j/4}(x_j)$. Apply Theorem 13.1, with a small loss $u > 0$, to get

$$D_h(x_j, x_{j+1}) \le H_u \min\{\delta_I, K_a R_j^a\}^{\alpha-u}.$$

The dyadic sum satisfies

$$\sum_{j\ge 0} \min\{\delta, K_a R_j^a\}^{\alpha-u} \le C\delta^{\alpha-u}(1 + \log(1/\delta) + \log K_a) \le C_u(1 + \log K_a)\delta^{\alpha-2u}, \qquad 0 < \delta < 1.$$

Let $j \to \infty$ and use continuity at p. Uniformly over all the cells and their points,

$$\operatorname{diam}_{D_h} C \le 2C_u H_u (1 + \log K_a)\delta_I^{\alpha-2u}.$$

The random prefactor has moments of every order. Choose $u < \varepsilon/8$, absorb the logarithmic factor in δ_I, and apply an arbitrarily high moment of that prefactor. Together with the filling and void-mass exceptions, this proves (13.10). □

The finite-range reference and permanent guards. The same result holds for the stationary finite-range field X used to define the anchor normalizer. To see this with quantitative moments, use the full-scale white-noise construction to couple $h = X + Z$, where h is pinned outside a sufficiently large inspection buffer and Z is smooth Gaussian there. On that buffer the reference Weyl rule defines $D_X = e^{-\xi Z} \cdot D_h$, and the common chaos convention gives $\mu_X = e^{-\gamma Z}\mu_h$. Normalized local shell costs and inverse normalized GMC masses change by factors bounded by $e^{C\|Z\|_\infty}$. The latter has moments of every order, by Gaussian supremum concentration. Hölder's inequality therefore transfers all the moments used above. The coarse covariance estimates for X have the same logarithmic coefficient. This is a moment transfer through an explicit smooth coupling, rather than through an arbitrary absolute-continuity density.

Furthermore, the event in Proposition 13.2 is valid simultaneously after insertion of any locally finite deterministic guard set in the construction buffer: every witness ball for a guarded cell is still empty of the original Poisson cloud. A fixed separated anchor connection must cross a fixed

reference shell. Its reference crossing cost has every negative moment by [8, Proposition 4.7], also after the preceding smooth transfer. A cover by N guarded cells can connect its ports at reference cost at most N times the maximum cell diameter. Consequently the finite anchor observable satisfies, for every $\varepsilon, B > 0$,

$$\mathbb{P}[F_I < I^{\alpha-\varepsilon}] \le C_{\varepsilon,B} I^{-B} \tag{13.13}$$

for sufficiently large I.

Consequence and limitation. Adjacent Voronoi cells touch, so a graph path of n edges in a fixed buffer has reference endpoint distance at most $(2n+2)$ times the maximum reference cell diameter. Thus the hop metric has a lower exponent at least ξ/γ, with the same estimate for local first-exit passages. This conclusion is stronger than Euclidean mesh control. It is still an exponent estimate: after division by an arbitrary $I^{\xi/\gamma+o(1)}$ normalizer, positivity of a limiting crossing does not follow merely from this proposition.

14 Metric-ball volume inputs and Poisson counts

Set $\Delta = \gamma/\xi$ and $\alpha_* = 1/\Delta$. The volume estimate below is a previously established input. The new deductions in this section count Poisson nuclei and follow reference paths through their Voronoi cells.

Proposition 14.1 (Established uniform upper metric-volume exponent)**.** *For the reference metric, every compact $K \subset \mathbb{R}^d$ and every $\eta > 0$, almost surely there is $\varepsilon_0 > 0$ such that*

$$\mu_h^\gamma(B_\varepsilon^D(x)) \le \varepsilon^{\Delta-\eta} \quad (x \in K,\ 0 < \varepsilon < \varepsilon_0).$$

This is the upper bound in (2.2), quoted from [2, Proposition 3.1] for $d \ge 3$ and [6, Theorem 1.1] for $d = 2$. The fixed weak-reference version follows by the deterministic-multiple observation in Section 2.

Proposition 14.2 (Uniform counts in small metric balls)**.** *Assume Proposition 14.1. Conditional on the field, let $\mathcal{P}_I$ be sampled from $I\mu_h^\gamma$. Fix a compact K, $\kappa > 0$, and put $R_I = I^{-\alpha_*+\kappa}$, with $\kappa < \alpha_*$. For every $\tau > 0$,*

$$\sup_{x\in K} \#(\mathcal{P}_I \cap B_{R_I}^D(x)) \le I^{\kappa\Delta+\tau}$$

with probability tending to one as $I \to \infty$. The statement is uniform over the field-dependent centers, including points on a reference geodesic chosen from the field.

Proof. Choose a deterministic polynomial Euclidean net fine enough that its D-covering radius is at most R_I, with probability tending to one; the local Hölder bound permits this. On the uniform volume event, each corresponding ball of radius $2R_I$ has conditional Poisson mean at most

$$I(2R_I)^{\Delta-\eta} \le I^{\kappa\Delta+(\alpha_*-\kappa)\eta+o(1)}.$$

Choose $\eta > 0$ so that $(\alpha_* - \kappa)\eta < \tau/3$. The Poisson upper-tail estimate, followed by the polynomial-size net union bound, gives the stated bound, with room to absorb the constant radius enlargement. Each radius-R_I ball centered in K is contained in one of these enlarged net balls. Independence between different counts is not used. □

Lemma 14.3 (Facet connectivity of the full trace-cell family)**.** *Let $S \subset \mathbb{R}^d$ be a nonempty locally finite set of distinct sites, and let $P \subset \mathbb{R}^d$ be compact and connected. Let $\mathcal{C}(P)$ consist of* all *closed Euclidean Voronoi cells which meet P. This family is finite, and its graph of codimension-one face adjacencies is connected. In particular, any two cells in this family can be joined using at most $|\mathcal{C}(P)| - 1$ facet steps. The assertion allows degenerate Voronoi vertices and deterministic guard sites.*

If every cell in $\mathcal{C}(P)$ has Euclidean diameter at most δ, the cells lie in the closed δ-neighborhood of P, and the straight nucleus-to-nucleus edges of such a route lie in the closed 2δ-neighborhood of P. Thus these cell traces and embedded edges are confined to any prescribed open neighborhood of P once δ is sufficiently small.

Proof. The nearest-site distance is bounded on P. A nucleus whose cell meets P therefore lies in one fixed bounded neighborhood of P, which contains finitely many sites. The ordinary intersection graph of the finite closed sets $C \cap P$, $C \in \mathcal{C}(P)$, is connected: otherwise their union would disconnect P.

Fix a point x shared by two such cells, and let S_x be its finite set of nearest sites. If S_x has more than one element, its sites lie on one sphere about x and each is a vertex of $Q_x = \operatorname{conv} S_x$. Indeed, the linear functional in the direction $p - x$ is uniquely maximized at p among these equal-radius sites. In a sufficiently small neighborhood of x, the remaining sites cannot compete. The cells incident to x are consequently the translated normal cones of Q_x. The vertex-edge graph of a polytope is connected, and each of its edges corresponds to a codimension-one common face of the associated normal cones in $\mathbb{R}^d$. This remains true when Q_x lies in a proper affine subspace, by including the orthogonal linear factor of the normal cones. Hence the cells containing x are connected by facet steps. They all meet P, so they already belong to $\mathcal{C}(P)$. Replace the edges of its intersection graph by these facet routes and erase loops to obtain the claimed bound.

For the final assertion, each cell contains a point of P, so its diameter bound places the whole cell within distance δ of P. If two successive cells have nuclei p, q and share a point w, then $|p - w|, |q - w| \le \delta$. The segment $[p, q]$ lies in $\overline{B}_\delta(w)$, while w is within δ of P. This proves the 2δ bound. The argument does not assert that the segment lies in the union of the two cells. □

Remark 14.4. Completeness of $\mathcal{C}(P)$ is essential here. An arbitrary selected subfamily whose union covers P need not be facet-connected at a degenerate contact point; the intermediate incident cells used in the normal-fan argument might have been omitted. In the upper hop estimate all cells meeting each reference-path piece are counted, so the lemma applies directly.

Proposition 14.5 (Sharp upper hop exponent from small cells)**.** *Assume the uniform reference-metric cell estimate*

$$\max_{C \cap K \neq \varnothing} \operatorname{diam}_D(C) \le I^{-\alpha_* + \varepsilon} \tag{14.1}$$

with probability tending to one, for every compact K and every $\varepsilon > 0$. Then, for each compact connected K contained in a connected open U, and every $\delta > 0$,

$$\sup_{x,y \in K} d_I(x, y; U) \le I^{\alpha_* + \delta}$$

with probability tending to one. Here paths in the graph are required to have their cell traces in U; an arbitrarily small enlargement handles the corresponding embedded-edge convention.

Proof. Choose a deterministic bounded connected open set V with $K \subset V \Subset U$. The internal metric of V is finite and continuous, and its diameter on the compact set K is finite. Thus there

is a finite random L such that every pair in K has a reference path of length at most L inside V. All such paths lie in the fixed compact buffer $K' = \overline{V} \Subset U$. One can restrict to events on which L is at most a fixed constant and then let that constant increase.

Set $R_I = I^{-\alpha_*+\kappa}$ with a small $\kappa > 0$, and apply (14.1) with exponent slack $\kappa/2$. Divide any of the reference paths into pieces of D-length at most R_I. There are $O(1 + L/R_I)$ pieces when the path has length at most L. Every cell meeting a piece has its nucleus within D-distance $2R_I$ of the initial point of that piece. Indeed, the contact point is within R_I and the nucleus belongs to the same cell, whose D-diameter is smaller than R_I.

Proposition 14.2 bounds the number of such cells by $I^{\kappa\Delta+\tau}$, uniformly over all pieces and all selected reference paths. The complete family of cells meeting a piece is facet-connected by Lemma 14.3, and hence gives a Delaunay route using at most its cardinality minus one edges. This also handles degenerate contacts or deterministic guards; an arbitrary selected subcover would not suffice for this argument. Concatenate the routes for consecutive pieces. Consequently

$$d_I(x, y; U) \le C(1 + L/R_I)I^{\kappa\Delta+\tau} \le C(1 + L)I^{\alpha_*+\kappa(\Delta-1)+\tau}.$$

Choose κ, τ sufficiently small that the exponent is below $\alpha_* + \delta$. The strict slack absorbs the random finite L on events of probability tending to one. The Euclidean mesh tends to zero, so cells and embedded edges meeting the fixed compact buffer eventually remain in U. □

15 An almost sure hop exponent, uniformly over real intensities

The preceding cell estimates and the cited reference volume bound give an almost sure conclusion over all real intensities. This step concerns the marked Poisson coupling and does not use rigidity of a candidate graph limit.

Use the notation

$$\alpha = \xi/\gamma > 0, \qquad \Delta = \gamma/\xi = 1/\alpha.$$

Let h be the subcritical log-correlated field, let $\mu = \mu_h^\gamma$, and let $D = D_h$ be a field-measurable weak reference metric with the sharp scale constants and shell estimates used above. In particular, the proved cell estimates and cited metric inputs give the following almost sure statements on every fixed compact buffer:

1. D is a continuous length metric inducing the Euclidean topology, and, for some deterministic $\sigma > 0$ and a finite random C, $D(x, y) \le C|x - y|^\sigma$ on that buffer.

2. For every $u \in (0, \alpha)$ and every fixed $\theta \in (0, 1)$, there is a finite H_u such that

$$\operatorname{diam}_D B_{\theta r}(z) \le H_u \mu(B_r(z))^{\alpha-u} \quad \text{if } r \le r_*, \quad \mu(B_r(z)) \le 1.$$

   This is Theorem 13.1. Also, for some $a > 0$ and finite $K_a \ge 1$, $\mu(B_r(z)) \le K_a r^a$ on the buffer.

3. For every $\eta > 0$, uniformly over centers in the buffer,

$$\mu(B_r^D(z)) \le r^{\Delta-\eta} \quad \text{for every sufficiently small } r > 0.$$

   This is the previously established upper volume bound quoted in Proposition 14.1.

The chaos is almost surely a full-support atomless Radon measure. For this corollary only finiteness, rather than moments, of the random constants in these almost sure statements is needed.

Conditional on h, construct one Poisson process

$$\Pi \subset \mathbb{R}^d \times (0,\infty), \qquad \text{intensity } \mu(dx)\,dt, \qquad \mathcal{P}_I = \{x : (x,t) \in \Pi,\ t \le I\}.$$

Let d_I be unit-edge distance in the facet-adjacency graph of the closed Euclidean Voronoi cells of $\mathcal{P}_I$. At a point on a cell boundary, any containing cell may be chosen. For an open U, write $d_I(\cdot,\cdot;U)$ for distance when every cell in the route is required to lie in U. The conclusions also hold if confinement is instead imposed on the straight nucleus-to-nucleus edges.

Corollary 15.1 (Uniform almost sure exponent)**.** *On an event of probability one, simultaneously for every compact $K \subset \mathbb{R}^d$, every connected open set $U \supset K$, every $\delta > 0$, and every $\varepsilon > 0$, for all sufficiently large real I one has*

$$I^{\alpha-\varepsilon} \le d_I(x,y) \le d_I(x,y;U) \le I^{\alpha+\varepsilon} \qquad (x,y \in K,\ |x-y| \ge \delta). \tag{15.1}$$

The upper bound holds simultaneously for all $x, y \in K$, without the separation condition. Both bounds hold for every choice of endpoint cells. Thus, whenever the separated pair set is nonempty,

$$\lim_{I\to\infty} \sup_{\substack{x,y\in K\\ |x-y|\ge\delta}} \left| \frac{\log d_I(x,y)}{\log I} - \alpha \right| = 0 \quad \textit{almost surely,}$$

and the same is true for the confined distances.

Proof. Put $L_n = 2^n$ and work on each slab $L_n \le I \le 2L_n$. All geometric tests of emptiness will use the lower cloud $\mathcal{P}_{L_n}$; all counting bounds will use the upper cloud $\mathcal{P}_{2L_n}$. The Voronoi cells themselves need not be nested as I increases.

Uniform Euclidean localization and empty masses. Fix an integer-radius Euclidean buffer. At any fixed positive mesh, finitely many deterministic open balls cover its enlargement, and each has positive μ-mass. Under the marked coupling each such ball therefore contains a point of $\mathcal{P}_{L_n}$ eventually, almost surely. Taking a countable collection of meshes and buffers, the maximal nearest-site distance of the lower clouds tends to zero uniformly on every compact set.

This filling assertion controls whole cells of *every* locally finite superset $S \supset \mathcal{P}_{L_n}$, and hence every current cloud in the slab. Indeed, suppose that every point in B_{R+1} is within Euclidean distance $r < 1/8$ of the lower cloud and that a cell C of S meets $\overline{B}_R$. Its nucleus p is within distance r of a contact point in $\overline{B}_R$. If C reached distance greater than $2r$ from p, convexity would give a point $z \in C$ with $|z-p| = 2r$ and $z \in B_{R+1}$. A lower-cloud point within distance r of z would contradict the defining nearest-site inequality. Consequently $C \subset \overline{B}_{2r}(p)$ and $\operatorname{diam}_{\mathrm{Euc}} C \le 4r$. In particular, whole cells meeting a compact set have Euclidean diameter tending to zero, uniformly for all real intensities in each slab. Their radial empty-ball witnesses lie in one fixed larger buffer.

We also need a mass bound uniform over empty Euclidean balls. For a fixed bounded buffer V, write $M = \mu(V) < \infty$. The Poisson VC empty-set estimate for Euclidean balls gives deterministic $c, C, v > 0$, depending only on dimension, such that, when $Lt \ge C$,

$$\mathbb{P}\left[\exists B \subset V :\ \mathcal{P}_L \cap B = \varnothing,\ \mu(B) > t \mid h\right] \le C(1+LM)^v e^{-cLt}. \tag{15.2}$$

One direct proof introduces an independent cloud of the same intensity. With conditional probability at least $1/2$, a selected empty witness of mass greater than t contains at least $Lt/2$ points of the second cloud. Conditional on the combined cloud, its labels are fair independent

coin flips; the number of different ball intersections with N points is at most $C(N+1)^v$. The probability that a set of at least $Lt/2$ combined points has only second-cloud labels is at most e^{-cLt}. Poisson moments of N give (15.2). It is enough to select from countably many rational balls: every open-ball witness has rational balls inside it with masses tending to its mass.

Choose a deterministic A large enough and set $m_n = An/L_n$. For almost every field, the right side of (15.2), with $L = L_n$ and $t = m_n$, is summable in n. Conditional Borel–Cantelli and then Fubini show that almost surely, eventually,

$$\mu(B) \le m_n \quad \text{for every } \mathcal{P}_{L_n}\text{-empty ball } B \subset V. \tag{15.3}$$

No high moment of $\mu(V)$ is needed for this almost sure statement.

A reference-metric cell mesh uniform across the slab. For clarity, apply the deterministic radial argument again, rather than infer a slab estimate from a fixed-intensity probability bound. Let C be a cell of any locally finite $S \supset \mathcal{P}_{L_n}$ meeting the prescribed compact buffer, let p be its nucleus, and let $x \in C \setminus \{p\}$. Put

$$x_j = p + 2^{-j}(x-p), \qquad R_j = |x_j - p|.$$

Convexity gives $x_j \in C$. The open ball $B_{R_j}(x_j)$ contains no point of S, and hence no point of the lower cloud. Also both x_j and x_{j+1} lie in $B_{3R_j/4}(x_j)$. The Euclidean localization just proved puts all these balls in V and ensures $R_j \le r_*$ for large n. The inner-ball estimate and (15.3) imply

$$D(x_j, x_{j+1}) \le H_u \min\{m_n, K_a R_j^a\}^{\alpha-u}.$$

For $0 < 2u < \alpha$, summing the geometric scales and using continuity at p gives

$$\operatorname{diam}_D C \le C_u H_u (1 + \log K_a) m_n^{\alpha-2u}.$$

The random prefactor is finite and independent of S and n. By taking u sufficiently small and absorbing the logarithm and this prefactor, it follows that, for every $e > 0$, almost surely eventually

$$\sup_{L_n \le I \le 2L_n} \max_{C \in \mathcal{V}_I:\ C \cap \overline{B}_R \ne \varnothing} \operatorname{diam}_D C \le L_n^{-\alpha+e}. \tag{15.4}$$

This proves the needed assertion for newly born cells as well as the cells present at the lower intensity. There is no monotonicity claim about individual cells or graph distances.

Uniform counts from the upper cloud. Fix $0 < \kappa < \alpha$, $\tau > 0$, and put $R_n = L_n^{-\alpha+\kappa}$. Choose $\eta > 0$ so small that $\eta < \Delta$ and $(\alpha - \kappa)\eta < \tau/3$. On a deterministic Euclidean grid of mesh $L_n^{-M_0}$ in a fixed enlargement of B_R, there are $O(L_n^{M_0 d})$ points. Choose the deterministic integer M_0 with $M_0\sigma > \alpha - \kappa$. The local Hölder estimate then makes this grid a D-net of radius less than R_n, eventually almost surely.

For almost every field, the uniform upper metric-volume estimate holds for all sufficiently large n. Every grid-centered ball of radius $4R_n$ then has upper-cloud conditional mean at most

$$2L_n(4R_n)^{\Delta-\eta} \le L_n^{\kappa\Delta+\tau/2}.$$

The conditional Poisson upper-tail bound at level $L_n^{\kappa\Delta+\tau}$, followed by the polynomial grid union bound, is summable in n. Conditional Borel–Cantelli therefore gives, almost surely eventually,

$$\sup_{x \in \overline{B}_R} \#\big(\mathcal{P}_{2L_n} \cap \overline{B}^D_{2R_n}(x)\big) \le L_n^{\kappa\Delta+\tau}. \tag{15.5}$$

Here a closed radius-$2R_n$ ball at an arbitrary center is contained in an open radius-$4R_n$ ball at a nearest grid point. The field estimates may begin at a field-dependent index; conditional Borel–Cantelli is applied only to the remaining summable tail. Neither independence of the different counts nor independence of the successive slabs is required. Since $\mathcal{P}_I \subset \mathcal{P}_{2L_n}$, this bound applies to every current cloud in the slab, including its newly born nuclei.

Confined upper bound for all pairs. Fix compact $K \subset U$, with U connected and open. Choose a bounded connected open V with $K \subset V \Subset U$. The internal metric $D(\cdot,\cdot;V)$ is finite and continuous. Indeed, a length metric inducing the Euclidean topology admits short internal connections in small neighborhoods, and the connected open set V is covered by such neighborhoods. Its diameter on K is therefore finite. Consequently there is a finite random L such that each pair in K can be joined inside V by a reference-rectifiable path of length at most L.

Use (15.4) with $e = \kappa/2$. Divide any one of these reference paths into at most $1 + L/R_n$ pieces of D-length at most R_n. The nucleus of every cell meeting a piece is within D-distance $2R_n$ of its initial point, since the cell's whole D-diameter is smaller than R_n. Equation (15.5) thus bounds the *complete* family of cells meeting each piece by $L_n^{\kappa\Delta+\tau}$. Lemma 14.3 connects this complete family by facet steps, even at degenerate contacts. Choose any containing cell at each partition endpoint and concatenate these routes. The result is valid for any choices of the original endpoint cells and satisfies

$$d_I(x,y;U) \le C(1 + L/R_n)L_n^{\kappa\Delta+\tau} \le C(1+L)L_n^{\alpha+\kappa(\Delta-1)+\tau}. \tag{15.6}$$

All whole cells and straight embedded edges in the routes lie in U for large n: their Euclidean distance from the path tends to zero, while $\overline{V}$ has a fixed positive distance from U^c. The count and mesh bounds already hold for every center and cell, so the argument is simultaneous over all pairs and reference paths. Choose κ, τ so small that $\kappa(\Delta+1)+\tau < \varepsilon/2$. The remaining exponent slack absorbs L and converts (15.6) into the asserted upper bound for every $I \in [L_n, 2L_n]$.

Global lower bound, including paths which leave the buffer. Choose R so that $K \subset B_{R-1}$. If the separated pair set is nonempty, continuity and positivity of D give

$$m_\delta := \min\left\{ \min_{\substack{x,y\in K\\ |x-y|\ge\delta}} D(x,y),\ D(K,\partial B_R) \right\} > 0.$$

Take any finite graph walk with N edges from a cell containing x to one containing y. Choose points on its successive common facets. Join consecutive chosen points, including x and y, by Euclidean segments in their common convex cell. This gives a continuous trace. No claim is made that these Euclidean segments have finite D-length.

If this trace does not hit ∂B_R, the triangle inequality applied to its successive chosen points gives

$$D(x,y) \le (N+1) \max_{C\cap\overline{B}_R\neq\varnothing} \operatorname{diam}_D C.$$

If it does hit ∂B_R, stop at its first such point z. All cells of this prefix meet $\overline{B}_R$, and the same finite triangle inequality bounds $D(x,z)$ by the right side. Either way,

$$(N+1)L_n^{-\alpha+e} \ge m_\delta$$

by (15.4). Choose $0 < e < \min\{\varepsilon/2, \alpha/2\}$. For all sufficiently large n, uniformly over the pair, the walk, and $I \in [L_n, 2L_n]$, this implies $N \geq I^{\alpha-\varepsilon}$. Taking the shortest unrestricted walk proves the lower bound. Such a finite walk always exists, for the full trace-cell lemma can be applied to the Euclidean segment joining the endpoints.

Finally take a countable intersection over integer-radius buffers, rational mesh sizes, and rational positive exponent losses in the estimates above. On this one event all geometric, volume, and count controls hold. Every compact K and open connected U can then be treated by the deterministic choices and pathwise arguments just given; no additional exceptional event depends on them. This proves the simultaneous formulation and its logarithmic consequence. □

Remark 15.2 (Guards and scope). The same conclusion holds after adding any one fixed locally finite set of distinct deterministic guard sites. The lower-cloud empty-ball and whole-cell arguments continue to hold. On the fixed buffer used for the upper routes there are only finitely many guards, so their contribution to the metric-ball count is absorbed by the strict exponent slack. The normal-fan trace-cell lemma already handles the resulting degeneracies. The corollary establishes the sharp almost sure power of the graph distance. It does not identify a normalized metric limit or prove that a subpower normalization converges.

16 Common local rules under changed field laws

The graph changes local Gaussian laws and conditions on coarse profiles. We must use the same Borel internal-metric rule for the actual reference in these experiments. The locality axiom in [14, Definition 1.1] holds for an LGF plus a possibly dependent random continuous function, as does its planar counterpart. The mixture argument below fixes the needed versions without changing the reference assignment.

Lemma 16.1 (Locality under mixtures of admissible inputs)**.** *Let $\mathcal{D}$ be one measurable weak exponential metric assignment, and let $(G_j)_{j\geq 1}$ be any countable family of input laws, each representable as an LGF plus a random continuous function. For every deterministic connected open set U, there is one Borel rule R_U on local generalized functions such that*

$$\mathcal{D}_{G_j}(\cdot,\cdot;U) = R_U(G_j|_U) \quad \textit{almost surely, for every } j.$$

The rules can be chosen simultaneously on a countable chart library and for a finite or countable family of actual weak assignments. The conclusion concerns the given assignments, not newly sampled metrics with matching marginal laws.

Proof. For each j, represent $G_j = H_j + f_j$, and absorb the additive constant of H_j into f_j so that all H_j have one fixed canonical pinned LGF law. On the disjoint union of these probability spaces choose a selector J with every $\mathbb{P}[J = j] > 0$. Then H_J still has that pinned LGF law. No independence between H_J and f_J is required. Hence $G = H_J + f_J$ is an admissible input for the primary locality axiom.

Apply locality to $\mathcal{D}_G$ on U and use the Doob–Dynkin factorization to obtain a Borel rule R_U. The fields and metric readouts can be encoded by countably many smooth test pairings and rational distance coordinates; thus the usual standard Borel form of that factorization applies. Conditioning the resulting almost-sure identity on $\{J = j\}$ gives the stated identity for the actual metric $\mathcal{D}_{G_j}$. Intersecting countably many probability-one events proves the final assertion. □

Corollary 16.2 (One fixed rule for every locally dominated input law)**.** *Fix one admissible input G_0 and a Borel local rule R^0_U for its actual internal metric. Let G be any other admissible input whose local field law satisfies*

$$\mathrm{Law}(G|_U) \ll \mathrm{Law}(G_0|_U).$$

Then this very same rule satisfies

$$\mathcal{D}_G(\cdot,\cdot;U) = R^0_U(G|_U) \quad \textit{almost surely.}$$

In particular the rule can be fixed before choosing a Gaussian coarse profile or a pin change, provided the local domination is proved for the resulting input law.

Proof. Apply Lemma 16.1 to the two laws G_0, G, obtaining a common rule R_U. It equals R^0_U almost everywhere for the local law of G_0. Domination gives the same equality for the local law of G, and the conclusion follows. This also shows why separate law-wise locality without the mixture axiom would have been insufficient. □

Conditional fine and coarse fields. The corollary has an adequate conditional form. Let T be a coarse environment and let $G(T, Z)$ be the actual field obtained from its fine noise Z. Suppose this joint experiment is an admissible input and, for almost every t, its conditional local field law is dominated by that of $G_0|_U$. The unconditional local law is then dominated as well. The corollary identifies the actual metric with R^0_U jointly, and disintegration gives the same identity for almost every coarse environment.

If a statement is needed for each fixed deterministic smooth profile t, apply the corollary to that particular admissible input law; the fixed rule R^0_U does not change. Joint measurability of the input and of the original assignment then permits integration over any profile law for which these conditional identities hold. This does not require an uncountable intersection of exceptional events, nor an assertion that every arbitrary continuous profile is a Cameron–Martin shift.

In the present application, rescaled fine noise with a fixed smooth coarse profile is locally an LGF plus a smooth function. On a strictly buffered chart it can be realized as a global admissible input by adding an independent large-scale LGF component and compensating it, together with the prescribed profile, by a random smooth function. A smooth cutoff supplies a global extension when only a smaller chart is inspected. The separate local Gaussian equivalence lemma proves domination on charts with an untouched pin portion outside the chart. The mixture argument supplies the metric-version identification after that Gaussian calculation; it does not replace it.

Scope. This gives the common local versions needed by Gaussian insertion, actual-array transport, and conditional coarse-profile experiments. It does not claim a single deterministic local rule on every generalized function or a continuous-shift invariant probability-one domain for all uncountably many inputs. Those stronger perfection claims are unnecessary for these countable arrays and locally dominated laws.

17 Upper-order cone records without a preassigned exponent

The record-scale selection follows the mechanism of [1, Lemma 14.4, p. 172]; the observable here is a Poisson port cost.

The following deterministic lemma removes an unnecessary exponent-existence assumption from the record construction. It does not prove regular variation.

Lemma 17.1 (Cone records at the upper growth order)**.** *Let $b : [1,\infty) \to (0,\infty)$ be continuous and locally bounded above and away from zero, and suppose*

$$0 < \alpha := \limsup_{I\to\infty} \frac{\log b(I)}{\log I} < \infty.$$

For every sequence $\varepsilon_n \downarrow 0$ there are $I_n \to \infty$ such that, simultaneously for all $J \geq 1$,

$$\frac{b(J)}{b(I_n)} \leq \max\{(J/I_n)^{\alpha-\varepsilon_n}, (J/I_n)^{\alpha+\varepsilon_n}\}.$$

Moreover $b(I_n) = I_n^{\alpha+o(1)}$.

Proof. Put $g(u) = \log b(e^u)$ and maximize $g(u) - \alpha u - \varepsilon|u - T|$ over $u \geq 0$. The expression tends to $-\infty$ as $u \to \infty$, so a maximizer exists. Choose $T \to \infty$ along upper-order peaks satisfying $g(T) - \alpha T \geq -\varepsilon T/4$. On every fixed interval $[0, R]$ the objective is at most $-\varepsilon T + O_{\varepsilon,R}(1)$, whereas its value at T is at least $-\varepsilon T/4$. Thus the maximizers escape every compact interval. For each ε_n, choose T_n so a maximizer u_n exceeds n. Maximality and the triangle inequality give

$$g(u) - g(u_n) \leq \alpha(u - u_n) + \varepsilon_n|u - u_n|, \qquad u \geq 0.$$

Exponentiate and put $I_n = e^{u_n}$. Taking $u = 0$ gives $g(u_n) \geq g(0) + (\alpha - \varepsilon_n)u_n$; the definition of the upper order gives the matching upper logarithmic bound. □

For the fractional-moment anchor normalizer $b(I) = (\mathbb{E}F_I^p)^{1/p}$, continuity follows from the exact local Cameron–Martin response identity. Indeed choose $p < q < 1$; the first-moment estimate bounds $\mathbb{E}F_I^q$, while the Gaussian shift density and all of its derivatives have uniformly bounded finite moments on compact shift intervals. Hölder's inequality therefore permits differentiation of $\mathbb{E}[R_c F_I^p]$. The weak polynomial mesh lower bound and the first-moment upper bound place the upper order in a compact interval $\alpha \in [c_0, 1]$.

Corollary 17.2 (Uniform fixed offsets at upper-order records)**.** *Use the records of Lemma 17.1 for this normalizer and set $J_n(v) = vI_n$, $A_n(v) = v^\alpha b(I_n)$ for fixed $v > 0$. Then, for $1 \leq J \leq J_n(v)$,*

$$b(J) \leq A_n(v) e^{\varepsilon_n|\log v|}(J/J_n(v))^{\alpha-\varepsilon_n}.$$

Consequently the past-envelope tail and upper compactness theorem applies at every fixed offset v, with constants in the limiting bounds independent of v. Any countable family of offsets can be retained jointly. In particular, for every $q < q_0$, the limiting qth moments of $F_{vI_n}/A_n(v)$ are uniformly bounded over the retained offsets.

Proof. Substitute $J = vI_n u$, $u \leq 1$, in the cone inequality. Eventually $\alpha - \varepsilon_n \geq \alpha/2$ and $A_n(v) \geq J_n(v)^{\alpha/2}$. For each fixed v the factor $e^{\varepsilon_n|\log v|}$ tends to one, so the same eventual envelope constant can be used before taking the limit at that fixed offset. Apply the past-envelope theorem and a countable diagonal extraction. □

Remark 17.3 (The missing conclusion is substantive)*.* Even these global cones, exact Gaussian shift identities, and a known polynomial growth exponent do not imply regular variation. Fix $0 < p < 1$, $0 < \alpha < 1$, and a centered nondegenerate Gaussian Z. For small $\epsilon > 0$ set

$$F_I(Z) = 1 + e^{\alpha(\log I + \gamma Z)}[1 + \epsilon\cos(\omega(\log I + \gamma Z))]^{1/p}.$$

It satisfies $F_{tI}(Z) = F_I(Z + \gamma^{-1}\log t)$ and has all positive moments in its own units. For sufficiently small ϵ it is increasing in I. Gaussian tilting gives

$$(\mathbb{E}F_I^p)^{1/p} = CI^\alpha[1 + a\cos(\omega \log I + \theta)]^{1/p}(1 + o(1)), \qquad 0 < a < 1.$$

Along phase-maximizing subsequences the multiplier profile is $t^\alpha[(1 + a\cos(\omega\log t))/(1 + a)]^{1/p}$, which is not multiplicative. This is an obstruction to the proposed inference, not a counterexample to the Poisson–Voronoi convergence conjecture.

17.1 A quantile exponent upgrades to the moment normalizer

Lemma 17.4 (Noncircular quantile-to-moment upgrade)**.** *Assume the proved anchor-network past-envelope tail theorem, with tail exponent $q_0 > 0$. Let $m(I)$ be the fixed sufficiently high quantile required by that theorem. If, for some $\alpha_* > 0$,*

$$m(I) = I^{\alpha_* + o(1)}, \qquad I \to \infty,$$

and the low-intensity quantiles are bounded, then for every $0 < p < q_0$,

$$b_p(I) := (\mathbb{E}F_I^p)^{1/p} = I^{\alpha_* + o(1)}.$$

The same upper estimate holds uniformly at every intensity, rather than only at selected records.

Proof. Fix $\epsilon > 0$. Choose C_ϵ such that $v(J) := \max\{1, m(J)\} \le C_\epsilon J^{\alpha_* + \epsilon}$ for all $J \ge 1$. The high-quantile hypothesis for v is tautological. At an arbitrary large intensity I, apply the past-envelope theorem with

$$s = \alpha_* + \epsilon, \quad A = C_\epsilon I^{\alpha_* + \epsilon}, \quad K = 1.$$

The denominator condition holds with any fixed $a_0 < \alpha_*$. Consequently $\mathbb{P}(F_I > CxA) \le C'x^{-q_0}$ uniformly in I. Integration gives $b_p(I) \le C_{p,\epsilon} I^{\alpha_* + \epsilon}$. If the chosen quantile level is $\theta < 1$, then for every $y < m(I)$, $\mathbb{P}(F_I > y) > 1 - \theta$. Letting $y \uparrow m(I)$ in the moment bound gives $b_p(I) \ge (1 - \theta)^{1/p} m(I)$. Now let $\epsilon \downarrow 0$. □

Remark 17.5. This lemma requires a proved quantile exponent for the actual finite guarded port observable. Reference-metric volume estimates alone cannot be inserted as that hypothesis without a comparison to the Poisson geometry. Once such a comparison is established, no higher-moment bootstrap is needed: the tail theorem is applied to a deterministic power envelope of quantiles, not to the moment normalizer one is trying to estimate. Regular variation and metric identification remain separate conclusions.

18 Local transfer to a field pinned outside the buffer

Lemma 18.1 (Local absolute-continuity transfer)**.** *Use a radial compactly supported smooth kernel κ, normalized by $\int \kappa^2 = 1$, with $\int \kappa \ne 0$, and let X be its finite-range field with scale interval $(0, 1)$. Let h_ρ be the canonical whole-space logarithmic field pinned by a compactly supported probability measure ρ of finite logarithmic energy. If U is a bounded open buffer and $\operatorname{dist}(\overline{U}, \operatorname{supp}\rho) > 0$, the laws of $h_\rho|_U$ and $X|_U$ are equivalent. Use the same measurable GMC map for both fields, with the same deterministic Euclidean regularization and counterterm. Their joint laws with the nested Cox processes of intensity $\lambda\mu_Y^\gamma$, $\lambda > 0$, restricted to U, are then equivalent as well.*

In particular, if local events E_n, at deterministic intensities λ_n and with arbitrary deterministic denominators b_n, have probabilities tending to zero for X, their probabilities tend to zero for h_ρ, with the same λ_n and b_n.

Proof. Extend the kernel white noise to all scales. Radiality and $k(0) = 1$ imply, against mean-zero test functions, that its covariance is exactly $\log(1/|x-y|)$: the additive constant in $\int_0^R k((x-y)/s)\,ds/s = \log R - \log|x-y| + c_k + o(1)$ disappears. Choose $U \Subset V$ with V separated from the pin, and $0 < a < 1$ so small that $2ar_\kappa < \operatorname{dist}(V, \operatorname{supp}\rho)$, where $\operatorname{supp}\kappa \subset B_{r_\kappa}(0)$. Then, on V,

$$h_\rho = X_{<a} + Z, \qquad Z = -\langle X_{<a}, \rho\rangle + H_{\geq a,\rho}.$$

The fine pin average is well defined by the finite-energy assumption and uses white noise disjoint from that determining $X_{<a}|_V$. The pinned coarse term is independent of the fine noise and smooth; at large scales its zeroth-order variance is integrable because pinning cancels the constant mode. Thus $Z|_V$ is smooth and independent of $X_{<a}|_V$.

The Cameron–Martin space of $X_{<a}$ is $H^{d/2}$, with equivalent norms for each fixed $a > 0$. Every realization of $Z|_V$ has a smooth compactly supported extension agreeing with it near $\overline{U}$, obtained by a cutoff. Conditional on Z, Cameron–Martin equivalence therefore applies on U; mixing preserves equivalence. The independent smooth decomposition $X = X_{<a} + X_{[a,1]}$ gives the same equivalence for $X|_U$. Applying the common GMC map and then the common Cox-process kernel preserves absolute continuity in both directions. Finally, absolute continuity of probability measures implies that any sequence of events whose probabilities vanish under one law also vanishes under the other. □

Remark 18.2 (Patches meeting the pin). The separation assumption can be relaxed when a patch omits a positive portion of the pin with positive clearance. Choose a second pin τ outside an enlarged patch and write $Y = h_\tau|_U$, $A = \langle h_\tau, \rho\rangle$, so $h_\rho|_U = Y - A \cdot 1$. The constant function is a local Cameron–Martin direction for Y. Gaussian projection gives $Y = Y^\perp + Z \cdot 1$, where Z is a nondegenerate Gaussian independent of $Y^\perp$. A smooth global Cameron–Martin function supported near the omitted portion of ρ, disjoint from both U and the second pin, can have nonzero ρ-average. It changes A while leaving Y unchanged, so $\operatorname{Var}(A \mid Y) > 0$. Conditional on $Y^\perp$, both Z and $Z - A$ therefore have nondegenerate Gaussian laws with positive densities everywhere. They have the same $Y^\perp$ marginal. Their disintegrations are equivalent, proving equivalence of the two pinned restrictions. Combined with the lemma, this covers small patches meeting a usual spherical pin. A finite patch cover then transfers the local upper-modulus bounds on any given compact set.

Normalization and locality. For the stationary reference, conversion from white-noise GMC to the chosen Euclidean convention is one fixed multiplicative constant. Absorb it in the reference intensity *before* selecting record intensities, or rescale that intensity subsequence once while retaining its numerical denominators; no comparison between $b(I)$ and $b(cI)$ is asserted. Separately variance-normalizing the pinned field generally introduces a spatial deterministic density and is not the common-map convention of the lemma. Guarded local cell-trace events qualify directly. Unguarded graph observables require a filling/localization event confining all relevant cells to the buffer, with its vanishing failure probability added separately. This transfer gives no noncollapse conclusion.

19 Finite-projection localization of countable passage arrays

The localization assertion for a countable intensity-offset array has an order of quantifiers. At a fixed offset, a filling event localizes *every* walk record in the inspected compact chart simultaneously. For finitely many fixed offsets and charts, a finite union of failure events has probability tending to zero. No common finite-resolution filling event is asserted for all countably many offsets.

Indeed, let W be a bounded observation window and let $v_j \downarrow 0$. In the nested Cox coupling, conditional on the field, the first arrival time in W is strictly positive almost surely when $0 < \mu_h(W) < \infty$. For each finite I, some $v_j I$ therefore has no sampled point in W. A simultaneous filling event for all j has probability zero. The countable array is handled in its product topology.

Lemma 19.1 (Localization in a countable product)**.** *Let $\mathcal{A} = \prod_{j\geq 1} \mathcal{A}_j$, where each $\mathcal{A}_j$ is Polish, and let $A_n, \widehat{A}_n$ be coupled $\mathcal{A}$-valued random variables. Suppose for every finite m,*

$$\mathbb{P}[\pi_m A_n \neq \pi_m \widehat{A}_n] \longrightarrow 0,$$

where π_m retains the first m coordinates. Then $d_{\mathcal{A}}(A_n, \widehat{A}_n) \to 0$ in probability for a bounded metric generating the product topology. Consequently A_n and $\widehat{A}_n$ have the same joint subsequential limits with any background variables retained identically in the coupling.

Suppose, in addition, that $\widehat{A}_n$ is a local construction with a conditional kernel depending only on the field in its observation window. Every limiting local-kernel identity proved for its finite projections, including conditional product identities on separated windows, holds for the full countable limiting array.

Proof. Choose bounded compatible metrics $d_j \leq 1$ and set $d_{\mathcal{A}}(a, b) = \sum_{j\geq 1} 2^{-j} d_j(a_j, b_j)$. For any m, equality of the first m coordinates bounds this metric by 2^{-m}. Given $\varepsilon > 0$, first choose m with $2^{-m} < \varepsilon$ and then let $n \to \infty$. This proves convergence in probability. If B_n is retained identically, the same bound holds in a bounded product metric on (B_n, A_n) and $(B_n, \widehat{A}_n)$, so asymptotically equivalent sequences have the same joint weak limits whenever either is tight.

For the kernel assertion, first use each finite projection. Its localized and original limits coincide by the preceding argument. Conditional expectations of bounded cylinder functions therefore satisfy the claimed local identities. The cylinders form a countable generating π-system after choosing countable bases in the coordinate spaces. Their conditional identities extend to all bounded measurable functions by the conditional monotone-class theorem. The projected kernels are consistent because the limiting array is one joint random object; regular conditional probabilities on the Polish product give its full kernel. No simultaneous prelimit localization of the infinite array is used. □

Application to actual Poisson arrays. Enumerate the chart, intensity-offset, mark-number, and auxiliary record coordinates. For any fixed finite initial list, all its intensities $v_j I_n$ tend to infinity. The filling estimate on each buffer localizes its complete finite-walk record collection; taking the finite union gives the premise of Lemma 19.1. The localized construction may be defined for the entire countable list from the Poisson cloud restricted to the observation window. Its whole-array kernel is exactly local even at finite resolution. Only its comparison with the original whole-space array is asymptotic in the product topology.

If a new base offset v_k is selected before choosing its actual resolution $n(k)$, retain the first k relative offsets and charts. Choose $n(k)$ so late that every intensity in this finite list diverges, all its localization failures have total probability at most 2^{-k}, and all the separately required marked-law and calibration-moment approximations hold. This is possible because v_k and the finite relative list are fixed before $n(k)$ is chosen. No positive lower bound on the infinite list of offsets is needed.

Exterior restrictions in exact completion. Let V_1 be the preservation patch and let V_0 be the unbounded exterior patch, with positively separated closures, containing the compact normalization pin. Define the exterior passage restriction from its countable collection of records

compactly confined in V_0. For every finite collection of those records, choose a bounded open $W_m \Subset V_0$ containing their compact chart buffers and the pin. Apply finite-projection localization and the finite-window Gaussian kernel-closure argument to W_m and the interior buffer. Retaining the actual full white-noise background does not alter their local conditional kernels. Exhaust V_0 by such windows and use the conditional monotone-class theorem. The entire exterior record then has a kernel depending only on $h|_{V_0}$, and its conditional joint kernel with the interior restriction is the required product.

Thus a localized field shift which is zero on V_0 preserves the exterior kernel, including its entire countable offset array, at the *limiting* level used for exact completion. The proof requires neither a single finite-resolution filling event on the unbounded exterior nor a global Gaussian density for the full exterior field.

20 The lattice detour and full marked-path interface in every dimension

The compactness of path records is dimension independent. Its probabilistic inputs are the already established local upper joining modulus, compact internal-diameter tightness, and vanishing Euclidean mesh and normalized edge cost. The lattice routing and marked-path arguments below apply in every $d \geq 2$.

20.1 The boundary convention in the lattice routing step

Nearest-neighbor adjacency on $\mathbb{Z}^d$ means ℓ^1 distance one; star adjacency means ℓ^∞ distance one. These must not be interchanged. For example, the external nearest-neighbor vertex boundary of a single vertex in $\mathbb{Z}^2$ is not nearest-neighbor connected. The valid boundary statement needed here is the following.

Lemma 20.1 (A boundary facing one complementary component)**.** *Let $C \subset \mathbb{Z}^d$, $d \geq 2$, be finite and nearest-neighbor connected, and let V be a nearest-neighbor component of $\mathbb{Z}^d \setminus C$. Then*

$$B(C, V) = \{v \in V : \operatorname{dist}_1(v, C) = 1\}$$

is star-connected.

Proof. The external edge boundary of any finite nearest-neighbor connected set has a connected family of dual $(d-1)$-plaquettes, with adjacency through $(d-2)$-faces; see [13, Lemma 3.1]. If V is the unbounded complementary component, apply this fact to C and take the V-side endpoints of those edges. Adjacent dual plaquettes have V-side endpoints at ℓ^∞ distance at most one, giving the claim.

If V is bounded, apply the same fact to the finite connected set V. Every vertex of C lies in the unbounded component of $\mathbb{Z}^d \setminus V$: connectedness of C joins it to a vertex outside the coordinate bounding box of V, and a straight ray from an extreme such vertex reaches infinity without meeting V. Every edge from V to its complement ends in C, since V is a component of C^c. Thus all C–V boundary edges are external for V. Taking their V-side endpoints again proves star-connectedness. This uses a plaquette boundary fact valid in every dimension, not a planar dual path or a Jordan curve. □

Proposition 20.2 (Confined open-lattice detours)**.** *Declare all sites outside a finite macro box open. Suppose every closed star-component has ℓ^∞ diameter at most R. There is a unique infinite*

open nearest-neighbor cluster $\mathcal{O}$. If $x, y \in \mathcal{O}$ can be connected by a lattice path contained in a set U, they can be connected by an open nearest-neighbor path inside the ℓ^∞ neighborhood of U of radius $R+3$. For a Euclidean connected open U, replace it first by its fixed one-mesh enlargement to obtain the same conclusion with a harmless additional constant. Every lattice site is within distance $R+4$ of $\mathcal{O}$, after increasing this fixed constant if needed.

Proof. All sites outside the macro box belong to one nearest-neighbor cluster, since $d \geq 2$, and every infinite open cluster reaches that exterior. This proves existence and uniqueness.

Fix the finite prescribed lattice path $P \subset U$ from x to y, let E be its $(R+3)$-neighborhood, and let C be the open component of x using only sites in this finite set E. Suppose $y \notin C$, and let V be the complementary component of C containing y. A last exit from C along P gives a vertex $v \in B(C,V) \cap P$. There is also $w \in B(C,V) \setminus E$. If V is bounded, follow an infinite open path from y until its first entry into C; its preceding vertex lies in V, and it cannot belong to E, since an open neighbor in E would belong to C. If V is unbounded, take an infinite open path from x and its last exit from the finite set C; the first vertex after that exit lies in the unbounded component V and is again outside E.

Lemma 20.1 joins v to w by a star path in $B(C,V)$. Every vertex of this path inside E is closed: otherwise it would be an open neighbor of C inside E. Its portion up to the first exit from E therefore contains a closed star-connected set whose diameter exceeds R, a contradiction. The confined open connection follows.

For density, take a long coordinate-line segment through a given site, with both endpoints outside the macro box. Its endpoints belong to $\mathcal{O}$. The confined connection just proved stays within distance $R+3$ of that segment and crosses the perpendicular coordinate hyperplane through the given site. A vertex of that crossing is within distance $R+4$ of the site. □

At adjacent dyadic lattice levels, define a coarse vertex to be nice only if it is open at the coarse level and its entire fixed fine-grid neighborhood is open. Then every coarse nice edge can be replaced by its two fine edges, whose midpoint and endpoints are open. The infinite nice cluster is consequently contained in both spanning open clusters. Applying Proposition 20.2 to the nice configuration gives the required nearby common vertex. A simple confined route over distance $O(R)$ uses $O(R^d)$ vertices. With $R = O(j^2)$ this is exactly the $O(j^{2d})$ routing bound used by the anchor construction. The cluster-tail estimate is a separate probabilistic input; this deterministic argument asks for no independence between adjacent levels.

20.2 All finite ordered path records

The marked-path design is adapted from [1, Section 5.1, pp. 28–30, and Proposition 15.1, pp. 174–175]. We verify capture and attainability for Euclidean Delaunay walks.

Let Γ_n be the original embedded Delaunay graph and let $B_n \to \infty$ be the cost denominator. Work on a countable library of compact rational polyhedral chart sets K and open rational buffers. The library contains a compact chart inside every open neighborhood of any given compact confinement. For $m \geq 1$ retain every finite edge walk in K, every choice of $m+1$ ordered marked vertices, and every set of upper cost bounds, as a record

$$(x_0, \ldots, x_m; R_1, \ldots, R_m; c_1, \ldots, c_m; c).$$

Here R_i is the nonempty compact polygonal range of the indicated subwalk, c_i bounds its number of edges divided by B_n, and $c \geq \sum_i c_i$. Repeated marks and constant subwalks are allowed. Take

the closure of these records in the compact metric space

$$K^{m+1} \times \mathcal{H}(K)^m \times [0,\infty]^{m+1}, \tag{20.1}$$

where $\mathcal{H}(K)$ is the Hausdorff space of nonempty compact sets and $[0,\infty]$ is compactified by $u \mapsto u/(1+u)$. Allow the empty closed collection as well. Its hyperspace is compact metrizable. The countable product over charts and numbers of marks is therefore compact metrizable.

These are measurable random closed collections. In a bounded observation window the cloud has a measurable enumeration, and the finite walks form a countable collection. Euclidean Delaunay adjacency is a Borel finite-configuration condition. It is enough to use rational upper cost bounds before taking closure. Hitting an open subset of (20.1) is then a countable union of measurable finite-walk tests. Local stabilization below extends the construction to the original infinite cloud.

Deleting an interior mark unions its two neighboring ranges and adds their costs. Both operations are continuous on the compactified spaces. Deleting an initial or terminal marked piece discards its range and subcost while leaving the total upper allowance unchanged. This is also continuous. Inserting a repeated mark is continuous as well. Their finite projection identities pass through Hausdorff limits. We do not identify an intersection with a closed chart boundary by a continuity assertion; all comparisons between different chart coordinates use strict open enlargements.

Theorem 20.3 (Marked compactness, attainability, and capture)**.** *Assume the following local estimates along the sequence: the Euclidean mesh tends to zero in probability on every compact; for each $K \Subset U$, nearby graph vertices in K admit paths inside U whose normalized costs tend uniformly to zero with their Euclidean separation; and normalized internal diameters of compact sets in connected open buffers are tight. Then the full record array has a joint subsequential extraction with the field, the actual reference metrics, Gaussian scores, and every retained fixed-offset experiment. In an almost surely convergent representation it has these properties:*

1. *Every finite-cost limiting record has actual discrete walk approximants with any prescribed positive loss in endpoints, range, and cost.*
2. *Every sequence of bounded-cost discrete walks in a fixed compact set has a passage record subsequence. Every fixed finite number of ordered subwalks can be captured simultaneously, with nonnegative subcosts whose sum is at most the total cost.*
3. *Endpoints converging to the same point can be joined with cost tending to zero in every fixed open margin. Finite lists of passage records concatenate, with arbitrarily small cost loss, in every open neighborhood of their ranges.*
4. *For connected open U, the infimum $F_U(x,y)$ of costs of passage records with endpoints x, y and compact range in U is a finite continuous pseudometric. Zero distances are permitted here.*
5. *For each finite cost cost bound and each $\delta > 0$, passage records admit ordered cost subdivisions with subcosts at most δ, up to an arbitrarily small allowance, using a bounded number of marks. All these assertions persist under further extractions with the same uniform estimates.*

No continuously parametrized limiting path and no lower metric bound are assumed.

Proof. The record spaces themselves are compact. We explain the extra endpoint coordinate, which is needed to join endpoints selected after observing the configuration. Fix $K \Subset U$ and a slightly larger compact set inside U. Choose a deterministic Euclidean grid of mesh $\delta_n \downarrow 0$ sufficiently

slowly. At each grid vertex choose its nearest cloud site. The mesh estimate and the upper joining estimate allow this choice so that the maximal joining cost over pairs at Euclidean distance $O(\delta_n)$ tends to zero in probability, and the query-to-site error is at most δ_n with probability tending to one. A diagonal choice works for the first k charts and errors $1/k$ in successive blocks of indices, and hence for the countable library.

Let d_n^U be the internal normalized graph distance, possibly infinite. Replace it for this auxiliary construction only by $d_n^U \wedge T_n$, where $T_n \to \infty$. This remains a finite pseudometric. Use a continuous partition of unity subordinate to the deterministic grid neighborhoods; it assigns to each $x \in K$ a probability measure $p_n(x)$ on the nearby selected cloud sites. The optimal transportation cost with ground pseudometric $d_n^U \wedge T_n$ defines a continuous pseudometric

$$\widetilde{d}_n^U(x,y) = W_{d_n^U \wedge T_n}(p_n(x), p_n(y)).$$

On the mesh event every support point is within $O(\delta_n)$ of its query point. The transportation triangle inequality bounds the uniform difference between this interpolation and the capped query distance by twice the maximal such short joining cost. Compact internal-diameter tightness removes the cap on K with probability tending to one. A zero auxiliary pseudometric can be used on the vanishing-probability event where the local construction is undefined.

The assumed upper modulus and diameter tightness give tightness of $\widetilde{d}_n^U$ in $C(K \times K)$. Retain these interpolations, their vanishing error coordinates, and the full untruncated path records jointly. All fixed background coordinates have tight laws in their usual separable realizations, so Prokhorov and a countable diagonal give joint extraction. Skorokhod representation gives almost sure convergence. The limiting interpolations are continuous pseudometrics. Consequently every pair of approaching endpoints, including configuration-dependent vertices, has joining cost tending to zero. Auxiliary caps need not obey a finite-resolution affine identity: they disappear in every controlled limit. Exact index identities concern the untruncated records.

Hausdorff convergence of a record collection supplies a nearby prelimit record for each limiting record. Approximate that record once more by an actual walk record. This proves attainability with open loss. In particular, if the limiting whole range lies in a compact subset of an open chart, Hausdorff convergence puts all sufficiently late realizing ranges in that chart. This is the open-margin chart compatibility used throughout; no intersection with a closed chart boundary is being passed through a limit. Conversely, compactness of (20.1) captures every bounded-cost sequence and its finite marks. Addition is continuous at finite total costs, so the subcost inequality survives. The ranges are connected compact sets containing their marked endpoints, and their limits retain these properties. They need not yet be path connected. Adaptive first and last cuts are always performed on the actual walks before capture. Replace a cut point on a polygonal edge by an adjacent vertex; the spatial error is at most the mesh and the cost error at most one normalized edge. A finite prescribed list of such cuts can then be captured together. The cuts need not be stopping times for any filtration.

For concatenation, realize a finite list of passage records and correct their approaching endpoint discrepancies by the short joins in a small ball contained in the prescribed open enlargement. Concatenate the actual walks and capture the resulting marked record. Letting the join radius decrease along a diagonal gives an arbitrary desired open loss and vanishing joining cost. Reversal is immediate. Constant walks and the mesh estimate give zero-cost diagonal passage records. Finite chains of local joins along an ordinary continuous curve in connected U give finite costs between all endpoints. Concatenation gives the triangle inequality, and local joining gives continuity in both endpoints. This proves the pseudometric assertion, without identifying a raw closed-boundary internal-distance coordinate.

Finally, on an actual walk of cost at most T, mark the first vertex after each successive cost increment δ. The overshoot is at most B_n^{-1} and there are at most $T/\delta + 2$ marks. If an original recorded subcost exceeds the corresponding actual walk length, append a constant parameter interval at its terminal vertex of exactly that excess. Subdivide each original piece in this padded parameter and retain all its original endpoints. The new subcosts for that piece sum to its original subcost, and repeated marks account for the constant intervals. The total allowance c is an independent coordinate: if $c > \sum_i c_i$, leave its unused allowance unallocated rather than adding it to an original subcost. Thus projection recovers every original range, subcost, and total upper allowance exactly. For m original pieces the number of new marks is bounded by $T/\delta + 2m + 2$. Additional repeated zero-cost terminal marks place the records in one fixed finite coordinate. Capture them and use $B_n^{-1} \to 0$. This proves cost subdivision. Every operation just used employs compact coordinates, open margins, and the same uniform modulus, so it survives another controlled extraction. □

20.3 Local kernels, comparison preservation, and completion

The conditional-law architecture follows [1, Proposition 15.5 and Lemma 15.7, pp. 186–190], with Gaussian windows replacing its sewing experiments.

For a compact chart K inside a buffered window W, fixed-radius filling in the buffer makes all Voronoi cells with nuclei in K bounded inside a smaller window $W' \Subset W$. No site outside W can affect those cells: on W' a nearby site of the restricted cloud is closer than every outside site, and a filled surrounding layer prevents those cells from escaping W'. Thus all adjacencies between vertices of a walk in K are determined by the cloud in W. The entire closed record collection in K agrees with its localized construction, not merely each chosen walk. The failure probability tends to zero by the uniform filling estimate.

For a countable offset and chart library this agreement is required only on each fixed finite projection, as proved in Lemma 19.1. For positively separated buffered windows these localized arrays have conditional product Poisson kernels depending only on the corresponding field restrictions. Include each fixed-offset array as one output; no independence between offsets in one patch is asserted. Proposition 23.5, including its augmented Gaussian version, therefore passes the local kernels jointly with the full fixed background and all actual reference readouts. Countably many chart restrictions are handled by finite lists followed by diagonal extraction. The exact completion after a localized shift is then the regular conditional-distribution construction in Proposition 23.8; it does not require that the limiting kernel be a point mass.

To make the index convention explicit, let $K(g)$ be the base local kernel, and define the fixed-offset experiment at intensity vI_n in units $v^\alpha B_n$. Canonical GMC constant-shift covariance gives, with $c_v = \gamma^{-1} \log v$,

$$K^{(v)}(g) = (S_{v^{-\alpha}})_{\#} K(g + c_v). \tag{20.2}$$

For a deterministic similarity $\phi(u) = z + ru$, the physical cloud pulled back by ϕ has intensity $I_n \mu_{h \circ \phi + Q \log r}$. Consequently its pulled-back records in units $v^\alpha B_n$ have kernel

$$K^{(v)}(h \circ \phi + Q \log r - c_v).$$

These are transported-experiment identities for conditional laws, obtained from exact finite Poisson identities and a fixed local Gaussian density. They impose no pathwise Weyl rule on an unidentified passage pseudometric. Similarities of closed chart coordinates use corresponding charts; arbitrary restrictions are always taken with open margins.

Lemma 20.4 (Positive lower comparisons are intrinsic)**.** *Suppose every finite-cost passage record has cost at least $cD_h(x,y)$, for one $c>0$ and a continuous reference length metric. Then every passage record with compact range in open U has cost at least $cD_h(x,y;U)$.*

Proof. Subdivide a passage record of cost a into subcosts at most δ. Every point in a subrange can also be retained as an extra ordered mark: choose approximating points on actual polygonal walks and then adjacent vertices, using the vanishing mesh and one-edge cost. Capture the extra marks. The lower comparison therefore bounds the reference distance from the subrange's initial point to each of its points by $(\delta+o(1))/c$. The whole range is compactly contained in U, so it has positive reference distance from U^c. For sufficiently small δ, reference near-minimizers between successive marked endpoints remain in U; their lengths can be arbitrarily close to their ambient reference distances. The sum of those distances is at most a/c by the subcost inequality. Concatenation gives $D_h(x,y;U)\le a/c$. □

Proposition 20.5 (Comparison preservation)**.** *Under the common estimates of Theorem 20.3, a numerical upper comparison in the open-port formulation passes through controlled marked extractions. A common lower comparison $c\ge 0$ passes as well. Their intrinsic local forms are preserved under fixed locally equivalent field laws and fixed indexed changes.*

Proof. For upper comparison choose smaller endpoint ports, a compact buffer inside the requested open domain, and a strict cost allowance. Reference paths with those margins give bounded-cost passage records with closed endpoint and range constraints. Existence of such a record is a closed hyperspace test, including the retained actual reference cost in its cost bound. If its compact confinement is random, exhaust it by a countable library and then a finite subfamily outside an arbitrarily small probability event; the reference law is fixed. Portmanteau passes the finite tests. Exhaust the smaller ports, buffers, and rational allowances to recover the same upper constant.

A strictly violating lower passage record in a new limit would have attainable approximants with convergent endpoints, compactly confined ranges, and almost the same cost. Local uniform convergence of the retained reference readouts would contradict the old lower bound. For intrinsic bounds use the actual internal reference metric in a slightly larger open buffer. When the old lower bound is positive, Lemma 20.4 supplies this local formulation; for $c=0$ the assertion is nonnegativity.

These intrinsic tests are measurable from the local field and local record kernel. Fixed equivalent changes preserve their probability-one assertions in both directions. In (20.2) use only the established reference identity $v^{-\alpha}D_{g+c_v}=D_g$ on the testing patch to transfer the same numerical constants. The preceding extraction argument then preserves them through controlled moving-index diagonals. An optimal constant can improve in a new limit; retaining a separate saturation witness is what excludes improvement when needed. □

Lemma 20.6 (Controlled reextraction before rigidity)**.** *Let $I_n\to\infty$ be a sequence of actual discrete experiments in units $B_n\to\infty$. Suppose that for every fixed multiplier v in the countable indexed library the full experiment at vI_n, in units $v^\alpha B_n$, has a joint marked limit $\nu^{(v)}$ with the interface of Theorem 20.3. Suppose also that the calibration powers converge at each fixed v through some order $q_*>4p$. Constants in these fixed-v moment bounds may depend on v. Assume that all $\nu^{(v)}$ have the same finite upper comparison C to the actual reference metric. No positive lower comparison or rigidity is assumed.*

Then the indexed limit laws have tight local endpoint moduli and compact internal diameters, uniformly in v, and their calibration coordinates $Z^{(v)}$ have uniformly bounded q_*th moments, provided the fixed reference port variable W_+ has a finite q_*th moment. For every deterministic

sequence of indices selected after *these fixed-index limits, one may extract it and realize that extraction by a diagonal of the actual discrete arrays. This diagonal retains the full marked interface, calibration uniform integrability at every order below q_*, the fixed extra offsets, and the fixed-density local kernel identities. It may also retain any specified finite witness coordinates at each stage.*

Proof. The upper half of the finite-port argument does not use rigidity. For a prescribed anchor pair, follow a reference near-minimizer in its open trace cube, use its available upper passage record, realize that passage record, and correct the endpoints by local joining. Convert its graph walk to a cell-contact trace. The vanishing cell mesh keeps this trace in the prescribed cube. Inserting the finite permanent guard set increases a distinct-cell cover by at most its cardinality. The normalized finite insertion and endpoint errors vanish. Summing the finitely many tests gives, for every fixed v,

$$Z^{(v)} \le CW_+ \quad \text{a.s.} \tag{20.3}$$

Here W_+ is the same finite sum of actual reference internal port distances under the same background law for every index. Thus its q_*th moment bounds those of all $Z^{(v)}$ uniformly.

The same upper comparison bounds the auxiliary interpolation limits, with a strict domain margin. Fix $K \Subset V \Subset U$. A reference near-minimizer in V gives a passage record compactly in V; its actual realizing walks eventually lie in U. Therefore any retained interpolation limit in U satisfies

$$\widetilde{d}^{U,(v)}(x,y) \le CD_h(x,y;V), \qquad x,y \in K.$$

For endpoints sufficiently close, use a smaller ball compactly inside V instead. The reference local modulus is a common random modulus with a fixed law; the triangle inequality transfers it to both arguments of every $\widetilde{d}^{U,(v)}$. A finite chain of such balls gives a common tight diameter bound on each compact connected trace geometry. Consequently these continuous-coordinate laws are tight, uniformly in v. The full closed-record coordinates are already compact, while the Gaussian background and actual reference readouts have their fixed tight laws. The bound (20.3) supplies tightness of the calibration coordinates and uniform integrability below q_*. The countable product has a jointly convergent subsequence along any chosen deterministic indices, including any countable library of fixed extra offsets. Common upper comparison for all those offsets is inherited from the exact fixed-index identity and the reference constant-shift rule.

We spell out the order in the actual-array diagonal. First fix the jth chosen index v_j and the first j fixed extra offsets. Their fixed-index discrete limits are already available, or are appended by a further subsequence of the original base sequence. Only now choose an original resolution $n(j)$ sufficiently large. Match the first j marked coordinates and the continuous interpolation laws to within $1/j$ in a metric for weak convergence, and make their mesh, one-edge, local filling, interpolation-error and finite-guard failure probabilities at most $1/j$. In addition require the *actual* q_*th calibration moments at these indices to differ from their limit moments by at most $1/j$. This last requirement is available from the assumed fixed-index moment convergence; it does not follow from weak convergence alone. It gives a uniform bound on the actual q_*th moments, hence uniform integrability of every smaller calibration power. Tight interpolation laws and vanishing query errors give the same adaptive endpoint joins for the selected actual arrays. The preceding theorem then supplies attainability, capture, concatenation and subdivisions in their new limit.

Each v_j is a fixed finite number before $n(j)$ is chosen. We may therefore impose simultaneously

$$\frac{|\log v_j|}{\log I_{n(j)}} \to 0, \qquad v_j I_{n(j)} \to \infty, \qquad v_j^{\alpha} B_{n(j)} \to \infty.$$

At cone records also require $\epsilon_{n(j)}|\log v_j| \to 0$. These conditions keep the physical mesh and one-edge cost vanishing. Whenever the original units have a positive polynomial lower bound, as they do in the present applications, the logarithmic condition preserves such a lower bound in the selected physical intensity. The anchor estimates at each fixed index, or the interpolation laws just bounded by the common upper comparison, justify every finite modulus requirement before taking this diagonal. In a backward array choose its selected integer index first, then choose the original resolution so late that the successful-run length exceeds that index and that its already bounded relative-unit ratios are approximated to the desired accuracy. The same argument applies.

Finally, at each newly selected base index retain the whole array of its fixed extra multipliers. Localize the first j offsets and charts before choosing the original resolution $n(j)$, and require their total filling failure probability to be at most 2^{-j}. This is possible because that finite list is fixed first. The localized whole array has the exact conditional product kernels on separated patches. Lemma 19.1 identifies its limiting law with that of the actual whole-space array in the countable product topology. The kernel-closure proposition applies with the same fixed Gaussian background laws. A fixed extra multiplier is transported by its own fixed local Gaussian density; no density bound uniform in $|\log v_j|$ is invoked. This passes all the fixed-index identities to the reextracted family, with the actual reference readouts retained. Compact witness coordinates can be included in the finite lists before each choice of $n(j)$; any asserted witness inequality is then passed by its separately proved strict-margin or closed-test argument. □

Here are the moment and quantifier inputs in the two applications. At upper-order cone bases, every *fixed* v inherits the anchor past envelope, with coefficient $\exp(\epsilon_n|\log v|) \to 1$. The anchor tail provides every moment order below q_0. A fixed cutoff Cameron–Martin change, when needed, preserves any strictly smaller order by Hölder; its constant need not be uniform in v. Choose $4p < q_* < q^{**} < q_0$, both below the available reference-port moment order. Fixed-index q^{**} bounds imply the q_*th moment convergence used in the lemma. The upper annular comparison is initialized from the cone estimates *before* applying the lemma, and transfers the same finite constant to all fixed indices. Thus the lemma supplies the further indexed extractions used to prove rigidity; it does not assume rigidity to get them.

For the backward arrays in propagation, the successful-run past envelope first supplies the common *own-unit* marked bounds. The relative-unit lemma is proved at fixed indices in those units. Only after it bounds their deterministic relative factors are the quantum-unit families reextracted. The annular argument then gives a common finite upper comparison, after which the present lemma supplies every additional extraction required by rigidity. No uniform profile moment at arbitrary moving indices is assumed at any of these stages.

Remark 20.7 (What is and is not a compactness obstruction). The attainability, capture, ordered marks, endpoint joins, local kernels, intrinsic comparison transport, and controlled diagonals above introduce no further dimension-dependent requirement. In particular they follow once the anchor estimates and mesh estimates have been proved for the original graph. A topologically forced intersection of two unrelated embedded graph paths does not follow in dimension $d \geq 3$ and is not asserted here. Any annular comparison or forcing argument must supply its own valid connector mechanism. That is the remaining geometric issue, rather than an unspecified failure of the closed-path record construction.

21 Transferring actual arrays between finite-range and pinned fields

The stationary finite-range field X is useful for the deterministic normalizer and the block estimates. Properness and the fixed-anchor forcing argument are used under a whole-space logarithmic field instead. The two whole-space Gaussian laws can be singular. The construction below therefore extracts actual discrete arrays under each law; it never evaluates an X-almost surely defined global kernel at a sample from a singular pinned law.

Choose a nonnegative smooth unit-mass pin ρ_0, with connected compact support and positive on the interior of that support. Its support is placed outside the fixed calibration and initial witness windows. Let h be the unit-logarithmic whole-space field with $(h, \rho_0) = 0$. The reference assignment uses the same canonical GMC renormalization for h and X. An admissible bounded window W has a nonzero part of ρ_0 outside $\overline{W}$. In particular, a window whose closure misses the pin is admissible. A countable atlas of small buffered balls is admissible and covers all of $\mathbb{R}^d$.

Lemma 21.1 (Local Gaussian equivalence with a fixed smooth pin)**.** *On every admissible bounded window the laws of $h|_W$ and $X|_W$ are mutually absolutely continuous. A bounded constant shift on that window can likewise be implemented by a Cameron–Martin change of the pinned field. For finitely many positively separated windows the local Gaussian joint-to-product density holds, with the actual pinned quotient in the one block containing the whole pin if there is such a block.*

Proof. Put $s = d/2$. For a test distribution ϕ supported in $\overline{W}$, its pinned covariance norm is the homogeneous H^{-s} norm of

$$T\phi = \phi - (\int \phi)\rho_0.$$

These zero-mass distributions have support in one fixed compact set, so their homogeneous and inhomogeneous covariance norms are equivalent. Choose a smooth test η supported outside $\overline{W}$ with $\int \eta\rho_0 = 1$. Pairing with η controls $|\int \phi|$ by the norm of $T\phi$, whence $\|\phi\|_{H^{-s}} \le C\|T\phi\|_{H^{-s}}$; the reverse bound is immediate. Thus the pinned and finite-range local covariance norms are equivalent. Their covariance difference is smooth, because both have the same logarithmic singularity and the pin is smooth. After normalization by these local norms the difference is Hilbert–Schmidt. There is no null direction on an admissible window. The Gaussian density argument of Lemma 23.1 proves mutual absolute continuity.

For a constant c on W, choose a smooth compactly supported function which equals c on a slightly larger window and compensate its pin average by another smooth function supported where ρ_0 is nonzero outside that window. The resulting function has zero pin average and finite critical Sobolev energy, and hence is a pinned Cameron–Martin shift with the desired local constant value.

For separated windows, if some pin mass remains outside their union, the same cutoff argument rules out a linear relation among their local Gaussian spaces. The normalized cross-covariance operators are smooth and Hilbert–Schmidt, so the usual density proof applies. If the union contains the whole pin, connectedness of its support and positive separation force that support to lie in one block. Use Lemma 23.3 for its one-dimensional null quotient. This connected-support convention rules out a deterministic pin relation split entirely between two separated blocks. □

One local reference rule for both laws. Lemma 16.1 and Corollary 16.2 supply one local readout for the actual assignment under the countable family of admissible input laws used here. Thus locality applies to an LGF plus a possibly dependent continuous function, with the same local

rule under the locally dominated Gaussian laws. In the common white-noise construction, $h - X$ is a random smooth function on every compact set. Let J be an independent Bernoulli variable and put $G = h + J(X - h)$. This is a logarithmic field plus a continuous perturbation. Locality for this one mixture law supplies a Borel rule $\mathcal{D}_U(G|_U)$ for its internal reference metric. On the events $J = 0$ and $J = 1$, that same rule gives the actual $D_h(\cdot, \cdot; U)$ and $D_X(\cdot, \cdot; U)$, respectively. Thus the local field density also transports the actual reference readouts; separate law-wise measurability alone would not justify that conclusion. Countable mixtures and a countable chart library make these versions consistent for every fixed field law and shift used below. This argument uses the stated reference locality for such continuous perturbations; it does not assert graph convergence under arbitrary perturbations.

Proposition 21.2 (Actual-array transfer and preservation of comparisons)**.** *Let $I_n \to \infty$ and $B_n \to \infty$ be deterministic. Suppose the actual Poisson–Voronoi arrays under X, in units B_n, satisfy the local mesh, joining, and compactness estimates of Theorem 20.3, jointly for the required fixed intensity offsets. Extract an X limit retaining the actual reference internal metrics. Then, after a further subsequence, the actual arrays under h have a full joint marked limit with the actual whole-space field, white-noise background, reference metrics, and calibration coordinate. The local laws on each admissible atlas window are the fixed-density transfers of the corresponding X local laws.*

Consequently every deterministic intrinsic upper or lower comparison for the X law holds for the h law, and conversely. The comparisons extend from the admissible atlas to all compact confinements. In particular, a rigidity theorem proved under the proper pinned reference law identifies the intrinsic X law as well, with exactly the same numerical scalar. No assertion that the whole-space X metric is proper is involved.

Proof. For each admissible observation window W, localize the graph and its compactly confined records to the Poisson cloud in W. Fix first a finite collection of intensity offsets and compact chart buffers. Their buffer-filling events have total failure probability tending to zero under either field law; on their intersection, every walk record in those fixed coordinates agrees with its localized counterpart. Lemma 19.1 then gives asymptotic equivalence of the countable arrays in their product topology. No common finite-resolution filling event for all offsets is used. Conditional on the local field $g|_W$, the localized array has one universal kernel $K_{n,W}(g, \cdot)$: it is the same Poisson construction with intensity $I_n \mu_g$, including its countable offset coordinates and finite Poisson-record augmentations. This localized kernel is exactly local; its comparison with the original array is the finite-projection limit just specified. Gaussian recipe scores are retained separately with the actual background, through the fixed-background kernel argument below. In particular the finite arrays are not defined by a limiting kernel on an exceptional set of fields.

Let $R_W(g)$ be the fixed density of $h|_W$ relative to $X|_W$. For an event A_n concerning such an array,

$$\mathbb{P}_h(A_n) \le L\mathbb{P}_X(A_n) + \mathbb{E}_X[R_W \mathbf{1}_{\{R_W > L\}}] + o(1).$$

First take $n \to \infty$ and then $L \to \infty$. This transfers local mesh convergence, the uniform joining modulus, and tightness. It also transfers convergence of each local joint law: approximate R_W in L^1 by bounded continuous cylinder functions of the field and pass those functions through weak convergence. The error is uniform in n because the field marginal is fixed. Actual local reference readouts can be included using the common local Borel rule obtained from the mixture argument above, and then by the fixed-background version of Proposition 23.5.

Cover a compact trace region by finitely many smaller admissible atlas charts. Their local joining bounds give finite connections in every fixed connected open buffer by a deterministic

finite chain of intermediate ports. They also give a common local endpoint modulus. Their tight finite connection costs give compact internal-diameter tightness. Thus Theorem 20.3 applies to the actual arrays under h on the entire countable chart library, not only to one localized marginal. Jointly retain the fixed whole-space h, its white noise and its actual reference metric coordinates. These background laws are tight. Countable diagonal extraction yields a full canonical-h array. All open-margin restriction, attainability, ordered-capture, and concatenation identities come from these same actual arrays and hence are mutually consistent. The local marginals are the fixed-density limits already identified; another arbitrary choice of a global X kernel plays no role.

The kernel-closure proof now applies under h itself. It retains the full background and shows that the local passage kernels depend only on the actual local field, not on extra white-noise information conditional on that field. Lemma 21.1 and, for an unbounded exterior containing the pin, bounded exhaustion and Lemma 23.3, give the product kernels needed for exact localized completion. The fixed-offset identities follow from the exact finite Poisson constant-shift identity and the same local density passage. Thus the canonical-h array has the full interface used in the forcing argument. For nested-scale probability transfer one uses only the genuinely annular record coordinates, with surrounding joins confined away from the central hole. The ball connections used for ordinary tightness above are not inserted into a separated-annulus experiment.

An intrinsic comparison in one smaller chart is a probability-one property of its actual local reference readouts and passage kernel. The positive density R_W transfers it in both directions. For the upper comparison, follow a compact reference curve through finitely many small charts, and concatenate its realizable follows using the endpoint modulus. This preserves the numerical upper factor exactly.

For a positive lower factor c, fix a compact passage record range inside an open domain U and cover it by interiors of finitely many compact inner sets K_a, with triple buffers $K_a \Subset V_a \Subset W_a \Subset U$. Each W_a is contained in an admissible observation window. Use the local comparison in W_a and put

$$a_* = \min_a D_h(K_a, \partial V_a; W_a) > 0.$$

On actual realizing walks, successively cut at the first exit from V_a after choosing an inner set K_a containing the current point. Round cut points to adjacent vertices before capture. The captured piece has range in $\overline{V}_a \Subset W_a$, so its boundary endpoints are strictly inside the chart where comparison is used. Every completed exit has limiting cost at least ca_*. There can be only finitely many such exits for a fixed total cost bound: otherwise, for any prescribed integer m, simultaneous marked capture would give m disjoint pieces of cost at least ca_*, a contradiction once mca_* exceeds that cost bound. Capture the resulting bounded list, including its final incomplete piece, and sum the local intrinsic lower comparisons. Reference paths with arbitrarily small allowances in the corresponding charts concatenate inside U, proving the lower bound by $cD_h(\cdot,\cdot;U)$. The factor-zero case is merely nonnegativity. This proves the global intrinsic comparison under h; the identical finite-chart argument works under X in the reverse direction.

The retained calibration has positive probability of being nonzero under h whenever it does under X, since its fixed observation window was chosen disjoint from the pin and its local density is positive. Apply the fixed-anchor rigidity theorem under h. Its identified intrinsic law transfers back through the local densities, then through the finite-chart argument, to the X limit with the same scalar C. This completes the construction and both comparison transfers. □

Remark 21.3 (Why the pinned reference is proper). Only the properness of the canonical logarithmic reference law and its bounded compactly supported shifts is used. For a smooth pin, its large-radius coarse averages satisfy $h_r(0) = o(\log r)$ almost surely along dyadic radii. The all-negative-moment

annular crossing estimate then makes remote crossing costs diverge, since $\xi Q > 0$. Consequently its metric balls are Euclidean bounded and, by continuity, compact. A bounded Weyl shift preserves this property by uniform metric comparison. The stationary finite-range auxiliary metric need not be addressed, nor is properness inferred for arbitrary unbounded continuous perturbations.

Corollary 21.4 (Finite-energy repinning, including the usual sphere pin)**.** *Let σ be a compactly supported probability measure of finite logarithmic energy, with connected support; assume $d \geq 2$. Set*

$$h^\sigma = h - (h, \sigma),$$

where h has the smooth pin used above and the pairing is its Gaussian L^2 pairing. Any identified locally uniform graph convergence under the smooth pin, with one deterministic normalizer, transfers to the actual Poisson arrays under h^σ with that same normalizer and target D_{h^σ}. The transfer is joint with the actual field and reference metric and hence gives convergence in probability. In particular this applies when σ is uniform probability on the unit sphere. No convergence estimate uniform in a random intensity multiplier is required.

Proof. Finite logarithmic energy on a compact set is equivalent to finiteness of its inhomogeneous $H^{-d/2}$ energy, since the corresponding Bessel kernel has a logarithmic singularity and a bounded remainder on that set. Equivalently, $\sigma - \rho_0$ has finite homogeneous energy. Thus $Z = (h, \sigma)$ is a finite Gaussian variable and $h^\sigma = h - Z$ is well defined.

Choose a countable buffered atlas of windows W with nonzero mass of both σ and ρ_0 outside $\overline{W}$. It may intersect the pin sphere; disjointness is unnecessary. For tests supported in such a window, the map

$$T_\sigma \phi = \phi - (\int \phi)\sigma$$

has covariance norm equivalent to $H^{-d/2}$, by exactly the cutoff argument in Lemma 21.1: choose a smooth test outside $\overline{W}$ whose σ integral is nonzero, control $\int \phi$, and use $\sigma \in H^{-d/2}$. The same is true for the smooth pin.

The difference of the two local covariance forms is bounded and of rank at most two in these normalized Gaussian spaces. Indeed, writing $m(\phi) = \int \phi$ and $L(\phi) = \mathrm{Cov}(Z, (h, \phi))$, it is

$$-m(\phi)L(\psi) - m(\psi)L(\phi) + \mathrm{Var}(Z)m(\phi)m(\psi).$$

The functional m is bounded by the local covariance norm, and L is bounded by Cauchy–Schwarz. Both local covariance forms are coercive, so the finite-rank normalized perturbation is positive and bounded away from zero. The finite-projection Gaussian density calculation therefore gives mutual absolute continuity. This argument does not call the pin potential smooth on the sphere.

Apply the actual-array construction of Proposition 21.2, replacing X there by the smooth-pinned law and using these local densities. The reference versions are common because locality can be applied to the single mixture $h - JZ$ for an independent Bernoulli J. Every subsequence of the actual h^σ arrays has a further full marked extraction whose local intrinsic records are the transferred identified reference records. Finite chart chains and positive exit charges glue these identifications on each compact trace region.

The simultaneous constant Weyl rule gives, on this coupling, $D_{h^\sigma} = e^{-\xi Z} D_h$. It is proper because D_h is proper and Z is finite. Apply the compact-trace escape argument to this proper reference metric to identify unrestricted distances, then use finite endpoint grids and local joining to obtain local uniform convergence. The actual reference coordinate is retained throughout, so subsequential identification on the equality graph implies convergence in probability under the

original h^σ coupling. The pointwise random Weyl factor was used only for reference properness and local metric versions, never to replace the deterministic-intensity arrays by arrays at a random intensity.

Finally, uniform sphere measure has finite logarithmic energy in every $d \geq 2$: near a point, its radial surface element is of order $u^{d-2}\,du$, and $\int_0^1 u^{d-2}\log(1/u)\,du < \infty$. Its support is connected. Hence it satisfies the stated conditions. □

22 Reference geodesic cores and local Gaussian perturbations

The rigidity argument uses a common consequence of the planar and higher-dimensional confluence theorems: every geodesic segment contains a segment of a geodesic with endpoints in a fixed countable set. We first match the cited reference results to this statement, then transfer saturation to one deterministic endpoint pair. No reference confluence proof is repeated.

22.1 A common deterministic-anchor property

Lemma 22.1 (A local escape test globalizes a seed)**.** *Let D be a proper geodesic metric on $\mathbb{R}^d$ inducing its Euclidean topology, and let O be open. Suppose $a, b \in O$ have a unique $D(\cdot,\cdot;O)$-geodesic Q, compactly contained in O, of length L. Suppose every continuous path from a to O^c pays at least $b_O > L$ before its first exit from O. Then Q is the unique ambient D-geodesic between a and b.*

The exit bound can be checked using only the internal metric on O: it suffices to exhibit a boundary layer inside O which every exit must cross and whose internal crossing cost is at least b_O.

Proof. The path Q is an ambient competitor of length L. Every ambient minimizer consequently has length at most L and cannot leave O. It is therefore an internal minimizer and equals Q. The same argument excludes any shorter ambient competitor, without requiring information about the metric outside O. □

Remark 22.2 (The cited higher-dimensional cores and their escape condition). For $d \geq 3$, use the constructive coverage in [3, Lemma 3.2, Proposition 5.1, Sections 5.7 and 7.2–7.5]. Its local code records deterministic seed endpoints, rational core fractions, compact containment and internal uniqueness. Sections 7.4–7.5 give simultaneous coverage of every nonconstant geodesic subsegment, with arbitrarily small inspection regions.

The escape condition is obtained within that same construction. Choose the seed of Lemma 3.2 with endpoints $q, r_k \to q$ in a fixed inspection domain O. Its length L tends to zero, whereas the internal first-exit cost $E_O(q)$ is positive. Choose it with $L < E_O(q)$. The countable local test is this scale-invariant inequality. Rational witnesses are chosen between L/s and $E_O(q)/s$, with the local metric normalizer s of the cited construction, so the test is invariant under additive repinning. A compact internal boundary layer witnesses the exit bound.

The pocket of Proposition 5.1 places the whole seed in the minimum Weyl plateau, of weight $a > 0$, and all weights are at least a; see Section 5.7, p. 21. The modified seed has length aL and every exit costs at least $aE_O(q) > aL$. The strengthened local test is therefore satisfied by the successful constructions used in the cited finite-family and coverage argument. Lemma 22.1 makes each such seed ambiently unique.

Lemma 22.3 (Deterministic-anchor cores in every dimension)**.** *For the reference metric in any integer dimension $d \geq 2$, there is a fixed countable set $\mathcal{A} \subset \mathbb{R}^d$ with the following property almost surely. Every nonconstant ambient geodesic segment J compactly contained in an open set W contains a nonconstant segment of the unique geodesic Q between two distinct points of $\mathcal{A}$, and Q is compactly contained in W. The assertion is simultaneous over the segments J and sets W.*

Proof. For $d \geq 3$, let $\mathcal{A}$ contain all deterministic endpoints in the countable catalogue of Remark 22.2. Choose a covered core on an interior subsegment of J with inspection domain compactly contained in W. Its entire seed lies in that domain and is ambiently unique by the escape test. This proves the assertion in these dimensions.

For $d = 2$, take $\mathcal{A} = \mathbb{Q}^2$. Work on the simultaneous event of [11, Proposition 6], together with uniqueness for the countably many deterministic anchor pairs from [5, Theorem 1.2]. Parametrize J by reference arclength. Choose interior times $s < t$ so that, with $r = D_h(J(s), W^c) > 0$, one has $t - s < r/10$; if W^c is empty, choose any positive r for which the corresponding ball is contained in W.

Fix $0 < \epsilon < (t - s)/4$. Proposition 6 of [11] supplies $\delta > 0$ such that every geodesic between points within reference distance δ of $J(s)$ and $J(t)$ contains $J([s + \epsilon, t - \epsilon])$. Decrease δ so that $\delta < r/10$ and the two metric balls are disjoint. The reference metric has the Euclidean topology, so choose distinct anchors $a_0, a_1 \in \mathcal{A}$ in these balls. Then

$$D_h(a_0, a_1) \leq t - s + 2\delta < 3r/10, \qquad D_h(a_0, W^c) \geq r - \delta > 9r/10\,.$$

Every geodesic between the anchors stays in W, since a first exit would cost more than its entire length. Its compact image is therefore compactly contained in W. Their unique geodesic contains the stated positive-length core by Proposition 6. The shared core is a consecutive subarc of that geodesic, up to reversal. These deductions hold on one event for every choice of J and W, as required. □

22.2 Saturation at deterministic anchors

Proposition 22.4 (Saturation transfers to a deterministic anchor pair)**.** *Let D be a random proper geodesic metric inducing the Euclidean topology and satisfying the countable deterministic-anchor property of Lemma 22.3. For open domains U, let F_U be passage pseudometrics satisfying domain monotonicity, concatenation, and*

$$F_U(x, y) \leq C D(x, y; U)$$

for one deterministic $C < \infty$. Suppose, with positive probability, a nonconstant $D(\cdot, \cdot; W)$-geodesic P of length L, compactly contained in a fixed open W, is saturated:

$$F_W(P(0), P(L)) = CL\,.$$

Then there are deterministic distinct anchors a_0, a_1, rational $0 < t_0 < t_1 < 1$, and a deterministic bounded open set $U \Subset W$ such that, with positive probability, the unique ambient geodesic Q from a_0 to a_1, parametrized by fractional arclength, satisfies

$$Q([t_0, t_1]) \Subset U, \qquad F_U(Q(t_0), Q(t_1)) = C(t_1 - t_0) D(a_0, a_1)\,.$$

The whole anchor geodesic Q can be kept inside W.

Proof. Every subsegment of P is saturated in W. Otherwise a cheaper connection for that subsegment, concatenated with upper-comparison connections for the two remaining pieces, would contradict saturation of P. Choose a sufficiently short nonconstant interior subsegment whose length is less than its starting point's ambient distance to W^c. Any shorter ambient competitor must stay in W, so this subsegment is also an ambient geodesic.

The anchor property supplies a unique anchor geodesic $Q \Subset W$ sharing a nonconstant segment with it. Choose rational fractional times $t_0 < t_1$ strictly inside that common segment, and a finite union U of rational balls containing the chosen arc compactly, with $\overline{U} \subset W$. Domain monotonicity gives $F_U \geq F_W$ at its endpoints. The upper comparison in U gives the opposite bound at exactly the reference length $(t_1 - t_0)D(a_0, a_1)$, since the corresponding subarc of Q is an ambient geodesic contained in U. Thus equality holds.

The anchor pairs, rational times and rational-ball domains range over countable sets. Their events cover the positive-probability saturation event, so one deterministic choice has positive probability. Only the common anchor property is used here; local coding of the higher-dimensional seeds is not a hypothesis on the planar reference. □

22.3 Localized white-noise recipes in every dimension

For odd d, the Cameron–Martin precision operator of the log-correlated field is $(-\Delta)^{d/2}$, which is nonlocal. Consequently a compactly supported smooth field shift does not in general have a likelihood score determined by the field on the same shell. The planar argument's pairing against a local Laplacian cannot simply be relabelled as a critical Sobolev pairing. The following construction gives genuinely local scores after retaining a white-noise realization of the field.

Choose a nonnegative radial $k \in C_c^\infty(B(0,1))$ with $\int k = 1$, strictly positive on a neighborhood of the origin, and let W be white noise on $\mathbb{R}^d \times (0,\infty)$ with Lebesgue control measure. With a normalizing constant $c_k > 0$, define the field modulo constants by

$$h(x) = c_k \int_0^\infty \int_{\mathbb{R}^d} t^{-(d+1)/2} k((x-y)/t)\, W(dy\, dt). \tag{WN}$$

Use the compact-kernel white-noise representation in [8, Section 2.3, Lemma 2.9 and Proposition 1.10], whose scope includes every $d \geq 2$. Choose c_k so that the effective kernel $c_k k$ has squared L^2 norm one. The cited white-noise approximations converge against zero-integral test functions to the unit-log LGF, giving (WN) modulo constants. The fixed smooth average then chooses our representative. Whenever the coefficient c_k is suppressed in later white-noise formulas, it is absorbed into the kernel; the effective kernel then has $\int k^2 = 1$. The unit-integral convention in the local plateau recipe below keeps c_k explicit.

Lemma 22.5 (Plateau recipes with local Gaussian likelihoods)**.** *Fix $\varepsilon > 0$, a nonnegative smooth function η_ε supported in $(\varepsilon, 2\varepsilon)$ with integral one, and $g \in C_c^\infty(\mathbb{R}^d)$. Write*

$$K_t(x) = t^{-d} k(x/t), \qquad f_\varepsilon(x) = \int \eta_\varepsilon(t)(K_t * g)(x)\, dt,$$

and define the deterministic white-noise shift

$$u_\varepsilon(y,t) = c_k^{-1} \eta_\varepsilon(t) t^{-(d-1)/2} g(y).$$

Then the following assertions hold.

1. *Shifting W to $W + u_\varepsilon$ shifts h exactly to $h + f_\varepsilon$.*

2. *The support of f_ε lies in the 2ε-neighborhood of* $\operatorname{supp} g$. *If g equals a constant a throughout the 2ε-neighborhood of a set E, then $f_\varepsilon = a$ on E. Pointwise upper and lower bounds on the same neighborhood are likewise preserved.*

3. *The shift has finite energy*

$$\|u_\varepsilon\|_2^2 = c_k^{-2}\|g\|_2^2 \int \eta_\varepsilon(t)^2 t^{1-d}\,dt.$$

For $f_{z,r}(x) = f_\varepsilon((x-z)/r)$, the realizing shift is

$$u_{z,r}(y,t) = r^{-(d+1)/2} u_\varepsilon((y-z)/r, t/r), \qquad \|u_{z,r}\|_2 = \|u_\varepsilon\|_2.$$

4. *The Radon–Nikodym density of the shifted white-noise law is*

$$R_u(W) = \exp\{W(u) - \tfrac{1}{2}\|u\|_2^2\}.$$

Its score is measurable from W on the compact space–scale support of u. On a source event $|W(u)| \le L$, the pushforward of the restricted source law is dominated by $\exp\{L + \|u\|_2^2/2\}$ *times the unshifted white-noise law.*

Proof. Substitution into (WN) gives

$$c_k \int t^{-(d+1)/2} k((x-y)/t) u_\varepsilon(y,t)\,dy\,dt = \int \eta_\varepsilon(t)(K_t * g)(x)\,dt.$$

The support and plateau statements follow from $\operatorname{supp} K_t \subseteq B(0,t)$ and $\int K_t = 1$; nonnegativity gives preservation of inequalities. Direct integration proves the energy formula. The change of variables $(y,t) = (z + rv, rs)$ proves both the dilation formula and energy invariance.

The final assertion is the usual white-noise Cameron–Martin formula. If $W' = W + u$, then $W'(u) = W(u) + \|u\|_2^2$. On the source clipping event, the density at the output is consequently at most $\exp\{L + \|u\|_2^2/2\}$. Also $W(u)$ is centered Gaussian with variance $\|u\|_2^2$, so clipping has uniformly small failure probability for a fixed finite recipe family, with a bound uniform in the location and scale of any one translated and dilated copy of that family. □

Remark 22.6 (Scope and remaining record obligation). Buffered smooth walls, prescribed constant-shift patches, zero neighborhoods, and cheap plateaus can be constructed first for g and then smoothed at a sufficiently small deterministic ε. Strict geometric margins and slightly enlarged plateau neighborhoods retain all their required properties. A finite recipe catalogue has finite maximum energy, even though that bound can become large as its smoothing scales decrease. Subsequent physical dilation does not increase the energy.

This lemma repairs the *locality of the likelihood score in the enlarged Gaussian realization*; it does not say that the score is measurable from the field restriction $h|_O$ in odd dimension. To use it in the full convergence argument one must retain the relevant localized white-noise observations in the prelimit records, verify closure and conditional completion for those enlarged records, and verify the annular-transfer statements for the augmented local events. Those are additional obligations, not consequences of the energy calculation. The global Holder change-of-measure argument of [3], pp. 24–25, does not by itself establish these localized record facts.

23 Gaussian comparison and local conditional laws

Local conditional laws must survive extraction before a limiting metric has been identified. The fixed field marginal and a fixed Gaussian density provide that closure. The passage records may retain conditional randomness at this stage.

23.1 Gaussian comparison on separated local windows

We state the Gaussian input explicitly. Work either with the finite-range logarithmic field used in the screening construction, or with a whole-space logarithmic field pinned by a smooth unit-mass average supported outside the finitely many windows under discussion. Write $s = d/2$. On each fixed bounded window its covariance norm on test functions is equivalent to the H^{-s} norm. In the pinned case a test ϕ represents $\phi - (\int \phi)\rho_0$ in $\dot{H}^{-s}$, where ρ_0 is the exterior pin. The equivalence follows from a cutoff separating the window from ρ_0 and, at low frequencies, the cancellation of the total mass. Between windows with positively separated closures the covariance kernel is smooth.

Lemma 23.1 (Finite-window Gaussian joint-to-product density)**.** *Let $W_1, \dots, W_m$ be finitely many bounded windows with positively separated closures, and let $X_i = h|_{W_i}$ in the above experiment. Let μ_i be its Gaussian marginal law. The joint law μ is equivalent to $\mu_1 \otimes \cdots \otimes \mu_m$. Its density ρ belongs to $L^{1+\varepsilon}$ for some $\varepsilon > 0$ depending on the fixed windows. For $m = 2$ it belongs to L^2.*

Proof. Let $\mathcal{H}_i$ be the first Gaussian chaos generated by tests in W_i, equipped with its covariance norm. The cross-covariance operator $C_{ij} : \mathcal{H}_j \to \mathcal{H}_i$ is Hilbert–Schmidt: the cross kernel is smooth, and conjugating it by the local H^{-s} norm equivalences still gives a smoothing operator, of arbitrarily negative order on bounded windows. Consequently the full covariance operator R on $\bigoplus_i \mathcal{H}_i$, normalized to have identity diagonal blocks, satisfies $R - I \in \mathcal{S}_2$.

There is no nonzero linear relation among these Gaussian spaces. Indeed, such a relation represents a zero distribution of the form

$$\sum_i \phi_i - \Big(\sum_i \int \phi_i\Big)\rho_0, \qquad \operatorname{supp} \phi_i \subset \overline{W}_i.$$

Cutoffs separating the windows and the pin force each ϕ_i to vanish. The same argument without ρ_0 applies to the finite-range field, whose spectral covariance is strictly positive. Since $R - I$ is compact, injectivity now implies $R \geq aI$ for some $a > 0$. Also $R \leq bI$ for a finite b.

For a finite-dimensional compression, let λ_k be its covariance eigenvalues. The qth moment of its Gaussian density relative to the product standard Gaussian is

$$\prod_k \lambda_k^{(1-q)/2}\{q - (q-1)\lambda_k\}^{-1/2}.$$

Choose $q > 1$ close enough to one that $q-(q-1)b > 0$. The logarithm of each factor is $O_q((\lambda_k-1)^2)$ uniformly on $[a, b]$, since its value and first derivative vanish at one. The Hilbert–Schmidt bound therefore bounds these moments uniformly over finite projections. The density martingale converges in L^q and gives ρ. Applying the same argument to R^{-1} proves equivalence. With two blocks the eigenvalues occur as $1 \pm \sigma_k$, where the σ_k are the singular values of the cross operator, $\sup \sigma_k < 1$, and $\sum \sigma_k^2 < \infty$. The second moment is $\prod_k (1 - \sigma_k^2)^{-1} < \infty$. □

Remark 23.2. For three or more windows the preceding density need not be in L^2. For instance, three scalar Gaussian variables with pairwise correlation 0.8 have a covariance eigenvalue 2.6. An $L^{1+\varepsilon}$ density, and in fact merely a fixed L^1 density, is sufficient for the kernel argument that follows.

Lemma 23.3 (A normalization pin inside the exterior block)**.** *The conclusions of Lemma 23.1 also hold when the smooth normalization pin ρ_0 is supported entirely inside one bounded window W_0, positively separated from the other windows. In this case the exterior block's Gaussian test space is the quotient by its single null direction* $\operatorname{span}\{\rho_0\}$*. The joint-to-product density is taken relative to*

its actual, correspondingly pinned marginal law. Consequently local-kernel product closure applies to an interior patch and an unbounded positively separated exterior patch containing the pin, by bounded exhaustion of the exterior.

Proof. For a test distribution ϕ supported in W_0, set $T\phi = \phi - (\int \phi)\rho_0$. Its covariance norm is $\|T\phi\|_{\dot{H}^{-d/2}}$. On zero-mass distributions with support in a fixed bounded set this is equivalent to the inhomogeneous $H^{-d/2}$ norm. The map T is bounded in that norm, has kernel exactly $\operatorname{span}\{\rho_0\}$, and is a bounded projection onto the zero-mass representatives. Thus the covariance norm is equivalent to the quotient $H^{-d/2}$ norm.

The pinned covariance kernel is the logarithmic kernel plus smooth terms obtained by convolution with ρ_0. Smoothness here holds also on the pin support, because $\rho_0 \in C_c^\infty$. Between the separated windows its cross kernels are consequently smooth, and the normalized cross operators remain Hilbert–Schmidt after taking the quotient. If a joint linear relation existed, its distributional representative would be $T\phi_0 + \sum_{i\geq 1} T\phi_i = 0$. Restricting to each interior window gives $\phi_i = 0$; then $T\phi_0 = 0$, which is zero in the exterior quotient. The covariance is therefore strictly positive on the normalized quotient sum. The density proof is unchanged.

For an unbounded exterior patch, choose an increasing exhaustion by bounded windows containing the pin and retaining positive separation from the fixed interior patch. Apply the finite-window closure to each exhaustion member. The exterior record sigma-field is generated by their countable union, so the conditional product identities extend by the monotone-class theorem. No uniform bound on the density moments over that exhaustion is required. □

Lemma 23.4 (Fixed-background closure of local product kernels)**.** *Let $H = (X_1, \dots, X_m, Z)$ have a fixed law on a product of Polish spaces, and let μ_i be the law of X_i. Suppose that the joint law μ of $(X_1, \dots, X_m)$ is absolutely continuous with respect to $\mu_0 = \bigotimes_{i=1}^m \mu_i$. Let $A_{n,i}$ take values in Polish spaces and, conditional on H, have joint law*

$$\prod_{i=1}^m K_{n,i}(X_i, da_i).$$

If for every i,

$$\mathcal{L}(X_i, A_{n,i}) \Rightarrow \mu_i(dx_i)K_i(x_i, da_i),$$

then $(H, A_{n,1}, \dots, A_{n,m})$ converges in law to the measure with the original H marginal and conditional law $\prod_i K_i(X_i, da_i)$ given H. In fact, for every bounded measurable f of H and bounded continuous g_i of the ith record,

$$\mathbb{E}\left[f(H)\prod_i g_i(A_{n,i})\right] \longrightarrow \mathbb{E}\left[f(H)\prod_i \int g_i(a_i)K_i(X_i, da_i)\right]. \tag{23.1}$$

For every fixed density $L(H) \geq 0$ with $\mathbb{E}L(H) = 1$, the same convergence holds after changing the background law by L, and its conditional record kernels remain K_i.

Proof. Write $\rho = d\mu/d\mu_0$, and first assume $|g_i| \leq 1$. Set

$$q_{n,i}(x_i) = \int g_i(a_i)K_{n,i}(x_i, da_i), \qquad q_i(x_i) = \int g_i(a_i)K_i(x_i, da_i).$$

Both functions are bounded by one. For a bounded measurable f, choose a version of the conditional expectation and put

$$r_f(x) = \rho(x)\mathbb{E}[f(H) \mid (X_i)_i = x].$$

Define it to be zero where $\rho = 0$. It belongs to $L^1(\mu_0)$, and the conditional product hypothesis gives

$$\mathbb{E}\Big[f(H)\prod_i g_i(A_{n,i})\Big] = \int r_f(x)\prod_i q_{n,i}(x_i)\,\mu_0(dx).$$

Finite sums of functions $\prod_i \phi_i(x_i)$ with bounded continuous ϕ_i are dense in $L^1(\mu_0)$. Indeed, rectangle indicators span a dense subspace for the finite product probability measure, and each one-coordinate indicator can be approximated in $L^1(\mu_i)$ by continuous functions valued in $[0,1]$, using regularity of Borel probability measures on Polish spaces. A telescoping expansion gives the corresponding product approximation.

For $\varepsilon > 0$, choose $R(x) = \sum_{\ell=1}^{N}\prod_i \phi_{\ell,i}(x_i)$ with $\|r_f - R\|_{L^1(\mu_0)} < \varepsilon$. Replacing r_f by R in the displayed integral costs at most ε, uniformly in n. The integral of each term factors as

$$\prod_i \int \phi_{\ell,i}(x_i)g_i(a_i)\,\mathcal{L}(X_i, A_{n,i})(dx_i, da_i).$$

Marginal weak convergence sends this expression to

$$\prod_i \int \phi_{\ell,i}(x_i)g_i(a_i)\,\mu_i(dx_i)K_i(x_i, da_i).$$

Replacing R by r_f in this limit again costs at most ε. Letting $\varepsilon \downarrow 0$ proves (23.1); rescaling proves it for arbitrary bounded g_i.

The joint laws are tight because the law of H is fixed and every record marginal is tight. Products of bounded continuous tests of H and of the individual records determine probability measures on this finite product of Polish spaces. The preceding identity therefore identifies every subsequential limit with the stated conditional product measure, proving joint weak convergence.

For a fixed change of background law, first replace L by $L \wedge M$ in the tests. The proved identity applies with $f(H)(L(H) \wedge M)$ in place of $f(H)$. The discarded expectation is at most

$$\|f\|_\infty \prod_i \|g_i\|_\infty \mathbb{E}[(L - M)_+],$$

under both the prelimit and limiting law, since their background marginal is the same. It tends to zero as $M \to \infty$. This proves convergence of the product tests under the changed law. Tightness also transfers: for any event E,

$$\mathbb{E}[L\mathbf{1}_E] \le M\mathbb{P}(E) + \mathbb{E}[L\mathbf{1}_{\{L>M\}}],$$

and the second term is uniform in n. First choose M, then use the original joint tightness. The product tests therefore prove joint weak convergence under the changed background law. Its conditional record kernels are unchanged because the density depends only on H. □

Proposition 23.5 (Closure of local kernels under extraction)**.** *Let $W_1, \dots, W_m$ be bounded windows with pairwise positively separated closures, in a Gaussian setting of Lemma 23.1 or Lemma 23.3, and put $X_i = h|_{W_i}$. For each i, let $A_{n,i}$ be a random record in a Polish space. It may contain an entire countable array of passage costs, offsets, local metrics, and recipe marks. Suppose that, conditional on the full background field, these records have product law*

$$\prod_{i=1}^{m} K_{n,i}(X_i, da_i).$$

Suppose along one subsequence $(X_i, A_{n,i}) \Rightarrow \mu_i(dx_i)K_i(x_i, da_i)$ for every i. Then their joint limit with the field has conditional law $\prod_i K_i(X_i, da_i)$. The same assertion holds jointly with any fixed additional background random variable B whenever the prelimit conditional law given (h, B) is the displayed product. It also holds after any fixed equivalent change of the background law. Countably many windows, affine charts, constant offsets, and fixed changes of law can be included by diagonal extraction.

Proof. For finitely many windows, take $X_i = h|_{W_i}$ and place the entire fixed background field, together with any retained fixed background variable B, into Z in Lemma 23.4. Lemma 23.1 supplies the required joint-to-product absolute continuity of $(X_i)_i$. The conditional product hypothesis and the prescribed marginal limits are exactly the other hypotheses of the abstract lemma. It gives the claimed conditional law jointly with the entire background, and its final assertion handles every fixed equivalent change of that law.

An entire countable record array is a single Polish random element, so offsets and marks within one patch are retained together; they are not asserted to be conditionally independent of one another. For countably many windows, apply the finite-window result to every finite subfamily. The resulting identities determine the countable product conditional law by the monotone-class theorem. A diagonal subsequence includes the countable prescribed charts, offsets and fixed changes of law. All these statements concern the completed local records before any subsequent selection of witnesses. □

Fixed additional information. The joint law of the field and B must be fixed during extraction. One may take B to be the whole white noise; the proof uses $\mathbb{E}[g(B) \mid (X_i)_i]$ and does not invert the map from noise to field. Moving selected witnesses do not satisfy this hypothesis. For example, if H is uniform on $[0, 1]$ and $A_n = B_n = \mathbf{1}_{\{\{nH\}<1/2\}}$, they are conditionally independent given H, but their joint limit is (H, B, B) with an independent fair bit B. Selected witnesses are therefore discarded before the local product law is invoked.

23.2 Retaining local white-noise recipe scores in every dimension

Use the compact-kernel white-noise realization and the localized recipe directions constructed previously. For each i, retain a finite list $U_{i,a} = W(u_{i,a})$. Assume their time supports lie between ϵr_i and $C\epsilon r_i$, their space–time supports lie in the buffered ith patch, and the induced field shifts $f_{i,a}$ are supported in that patch. Distinct patches are chosen so these noise supports are disjoint and each $f_{i,a}$ vanishes on every other inspected patch and its normalization average. Recipe lists come from a fixed finite catalogue up to translation and dilation. Remove any zero or linearly dependent directions within each list.

Lemma 23.6 (Augmented Gaussian comparison)**.** *The finite-window density and local-kernel conclusions remain valid when X_i is augmented by the score vector $(U_{i,a})_a$. Moreover, consider normalized annular fields G_i on geometrically separated annuli, and assume their unaugmented covariance forms obey*

$$|\operatorname{Cov}(G_i(\phi), G_j(\psi))| \le C\theta^{|i-j|}\|\phi\|_{\mathcal{H}_i}\|\psi\|_{\mathcal{H}_j}, \qquad 0 < \theta < 1. \tag{23.2}$$

Here the field tests include subtraction of the local normalization average. The corresponding augmented Gaussian spaces satisfy the same estimate with a larger constant depending only on the fixed recipe catalogue. In particular, after taking a sufficiently sparse subsequence of annuli, for

arbitrary events E_i measurable in their augmented local data,

$$\mathbb{P}\Big(\bigcap_{i=1}^{N} E_i\Big) \le \prod_{i=1}^{N} \mathbb{P}(E_i)^{1/2}. \tag{23.3}$$

The same inequality holds when conditionally independent local kernel records are adjoined to the augmented Gaussian data.

Proof. The covariance of a score with a field test is the test applied to its induced shift. Thus, for $i \neq j$,

$$\operatorname{Cov}(U_{i,a}, G_j(\psi)) = 0, \qquad \operatorname{Cov}(U_{i,a}, U_{j,b}) = 0.$$

Only the original field–field cross form remains. The restriction of the field to white-noise times $t < \epsilon r_i/2$ is independent of every score in its own patch. On the unit rescaled patch its covariance is bounded below by a positive constant times the local $H^{-d/2}$ covariance norm. Indeed, with a fixed unit cutoff ε, its spectral density is

$$F_\varepsilon(\zeta) = \int_0^\varepsilon t^{d-1} |\widehat{k}(t\zeta)|^2 \, dt \ge c_\varepsilon (1 + |\zeta|^2)^{-d/2}.$$

Integrate over $0 < t < c\min(\varepsilon, |\zeta|^{-1})$ and use $\int k > 0$. The cutoff is below the recipe bands of its own rescaled block. Therefore

$$\|G_i(\phi) + \textstyle\sum_a c_a U_{i,a}\|_{L^2} \ge c\|\phi\|_{\mathcal{H}_i},$$

with $c > 0$ uniform over the fixed catalogue and its translates and dilates. This proves the asserted cross-covariance bound for the augmented spaces. It also shows that adjoining the finitely many scores preserves the Hilbert–Schmidt cross operators in the finite-window argument. A zero linear relation would have zero projection onto sufficiently small white-noise times, forcing its field-test distribution to vanish. Disjoint supports then force each field test to vanish, and the disjoint noise supports and quotiented Gram matrices force all remaining score directions to vanish. The positivity and density proof is consequently unchanged.

For the annular assertion, sparsify until the sum of the norms of the off-diagonal covariance blocks in each row is at most one. The full covariance is then at most twice its block-diagonal covariance. The Gaussian Holder inequality (equivalently, the Gaussian Brascamp–Lieb inequality) gives $\mathbb{E}\prod_i g_i \le \prod_i (\mathbb{E} g_i^2)^{1/2}$ for nonnegative local functions. Apply this to indicators and then pass from finite-dimensional Gaussian projections by conditional expectation. For kernel records, condition on the augmented Gaussian data, replace each indicator by its conditional probability $0 \le g_i \le 1$, and note that $\mathbb{E} g_i^2 \le \mathbb{E} g_i = \mathbb{P}(E_i)$. □

Remark 23.7 (What this resolves, and what it does not). The odd-dimensional nonlocality of $(-\Delta)^{d/2}$ prevents treating the field Cameron–Martin score for an arbitrary compactly supported shift as a local field observable. The localized white-noise recipes avoid that error: their scores are explicit local Gaussian variables, their density identities are exact, and the preceding lemmas retain these variables under extraction and annular probability comparison. The unaugmented annular estimate (23.2) is the separated zero-mass covariance bound of [14, Lemma A.1, equation (35)]. The new part above retains the local noise scores and the passage kernels. These facts establish the Gaussian/score portion of the local-kernel interface. They do not yet establish the additional transfer needed for labels selected from a whole coarse-average history; that joint shape/history comparison is a separate question. Nor do they supply a Weyl rule, noncollapse, or a rigidity conclusion for the passage limits. The separate shape/history argument below addresses the additional coarse-label transfer.

23.3 Exact completion after a localized noise shift

The order of completion and subsequent selection follows [1, Lemma 8.16, p. 79]; the augmented-noise identity needed here is proved below.

Proposition 23.8 (Augmented completion and bounded output density)**.** *Let W be the retained white noise, let $h = h(W)$ be its normalized logarithmic field, and let*

$$\mu(dW, dZ) = \mathcal{G}(dW)K(h(W), dZ)$$

be an extracted full passage-record law. Assume the compatible local restrictions have product kernels on positively separated patches, as established by Proposition 23.5. Let $u \in L^2(dy\,dt)$ be one deterministic localized recipe direction whose induced field shift f is zero on an exterior patch V_0 and equal to a constant b on a positively separated preservation patch V_1. Set $W' = W + u$ and $h' = h + f$.

Retain the old passage restriction on V_0, independently draw the V_1 restriction from $K_1(h|_{V_1}+b, \cdot)$, and complete all remaining coordinates using the conditional law of $K(h', \cdot)$ given those two restrictions. Before any recipe selection the output law is

$$\mathcal{G}_u(dW')K(h(W'), dZ').$$

If an arbitrary later selection event is contained in $\{W(u) \le L\}$, its output subprobability measure satisfies

$$\nu \le \exp\{L + \tfrac{1}{2}\|u\|_2^2\}\,\mu. \tag{SC}$$

The selection can use an old geodesic, auxiliary passage records, or information not measurable from the output.

Proof. The kernel-closure statement holds jointly with the entire fixed background W. Thus, even after retaining that background, passage records have conditional kernel $K(h(W), \cdot)$; no field measurability of the passage record itself is being assumed. Conditional on W', the retained exterior restriction has its correct target kernel because $f = 0$ on V_0, and the independently drawn interior restriction has its correct target kernel because $f = b$ on V_1. Product locality makes their pair exactly the target pair marginal. A regular conditional distribution of the full target record given that pair supplies the claimed completion. The unbounded exterior may contain the fixed normalization pin: Lemma 23.3, applied to bounded exterior exhaustions and then a conditional monotone-class limit, supplies precisely the product property needed in that case. The exterior records here are intrinsic records compactly confined to the exterior, not ambient distances whose minimizing paths might enter the changed region. Only the local preservation input is $h|_{V_1} + b$; the global shifted field is $h + f$ in the original pinned law class.

The Cameron–Martin density of the unselected output relative to μ is

$$\exp\{W'(u) - \tfrac{1}{2}\|u\|_2^2\}.$$

Since $W'(u) = W(u) + \|u\|_2^2$, this density is bounded by the constant in (SC) on the selected output. Restricting a source probability before pushing it forward can only reduce its output measure, which proves (SC). The argument is applied separately to each deterministic recipe in the finite catalogue; all auxiliary samples and completions are specified before selection. □

Remark 23.9. This proposition closes the augmented-record completion issue left open in the earlier localized-shift lemma. Its likelihood score is the local white-noise variable $W(u)$, including in odd dimension. No local representation of the nonlocal field pairing $\langle h, (-\Delta)^{d/2} f\rangle$ is asserted or needed.

24 Quantum-label annulus transfer without a Markov decomposition

The quantum-label transfer is adapted from [1, Lemma 8.12, pp. 72–75]. Gaussian covariance comparison supplies the all-dimensional replacement for the planar field decomposition.

This section supplies a higher-dimensional version of the all-point quantum-label transfer used in the restricted passage-law argument. The replacement for independent Dirichlet pieces is one comparison of the complete Gaussian shape and coarse-average vector with a product Gaussian experiment. No Brownian circle-average process is assumed.

Fix the testing annulus

$$S = \{1/16 < |x| < 4\} \subset \mathbb{R}^d,$$

and let ρ be uniform probability on the sphere ∂B_8. A smooth radial probability supported in a thin boundary layer of that sphere works equally well. The normalization sphere is outside the observed annulus. Let h be the whole-space logarithmic field, pinned by a smooth unit-mass average outside the physical compact under consideration. The field modulo constants has covariance kernel $-\log|x-y|$. Fix L large, put $r_j = e^{-t_j}$ with $t_j = t_0 + jL$, and suppose

$$|z_j - x| \le c_* r_j$$

for one point x and a sufficiently small fixed $c_* > 0$. Write

$$H_j = (h, \rho_{z_j,r_j}), \qquad G_j(u) = h(z_j + r_j u) - H_j \quad (u \in S), \qquad Z_j = H_j - H_{j-1} \quad (j \ge 1).$$

Here $\rho_{z,r}$ is the pushforward of ρ under $u \mapsto z + ru$. Every G_j has the same unit annular Gaussian law. Its test functions represent zero-mass distributions supported on $S \cup \partial B_8$.

Lemma 24.1 (Joint comparison of shapes and coarse history)**.** *Let M be the number of increments. Suppose $t_0 \ge aN$, $M \le C_0 N$, and all centres lie in a fixed compact set. For L sufficiently large and then N sufficiently large, the joint law of*

$$(G_0, \dots, G_M, H_0, Z_1, \dots, Z_M)$$

has a density $R_{N,M}$ relative to the product of its individual Gaussian marginals. The product law makes all the G_j, H_0, and Z_j independent. There is a constant C independent of M, N and the chosen ancestry such that

$$\mathbb{E}_{\mathrm{prod}} R_{N,M}^2 \le \exp\{CM/L + C\}. \tag{24.1}$$

The constant is uniform over physical compacts after increasing their starting value of N. Finitely many localized white-noise recipe scores may be included in each shape block, with constants depending on the fixed recipe catalogue, provided all induced recipe shifts vanish on the coarse normalization spheres and on other inspected annuli. More precisely, take a finite list of unit directions u_a supported in one buffered annular space–time region, with times in $[\epsilon_0, \epsilon_1]$ for fixed $0 < \epsilon_0 < \epsilon_1 < \infty$. Their induced shifts must be supported in a fixed annulus strictly between the central hole and the radius-8 pin. At scale j retain the variables $W(u_{j,a})$, where

$$u_{j,a}(y,t) = r_j^{-(d+1)/2} u_a((y - z_j)/r_j, t/r_j).$$

Choose the separation of scales so the supports of these directions are disjoint for distinct j, and every induced shift vanishes at every inspected coarse pin. Linear dependencies within each finite score list are quotiented out. The local passage and auxiliary kernels depend on the actual field in their observation patch, while the scores are retained as additional Gaussian background variables; they are not asserted to be functions of the local field restriction.

Proof. Set $\vartheta = e^{-L}$. Normalize each scalar Z_j by its standard deviation and H_0 by its standard deviation. On the first chaos of each shape use its covariance norm. This norm is uniformly equivalent to the $H^{-d/2}$ norm of its test on S: the exterior pin at radius 8 supplies the zero-mass representative, and a fixed cutoff separates that pin from S.

For the low-frequency part, write $T\phi = \phi - \langle \phi, 1\rangle \rho$. On the unit support, pairing against $e^{-i\zeta\cdot x} - \int e^{-i\zeta\cdot y}\rho(dy)$ gives

$$|\widehat{T\phi}(\zeta)| \le C|\zeta|\|\phi\|_{H^{-d/2}}, \qquad |\zeta| \le 1.$$

The integral of $|\zeta|^{2-d}$ near zero is finite. The spherical pin has finite $H^{-d/2}$ norm for $d \ge 2$, and a separating cutoff gives the reverse norm comparison. To quantify a smooth cross kernel, insert compact cutoffs and use a fixed torus. Its matrix in normalized $H^{-d/2}$ Fourier bases consists of its Fourier coefficients multiplied by $\langle k\rangle^{d/2}\langle l\rangle^{d/2}$. Integration by parts to an order greater than $2d+2$ bounds the squared sum by a constant times a sufficiently high squared derivative bound of the kernel. Thus the claimed smoothing estimates are Hilbert–Schmidt estimates also in odd dimension.

We first record the covariance estimates, including their summable off-diagonal structure. The constants below depend only on the fixed unit geometry, except for smooth exterior-pin errors that vanish as t_0 tends to infinity:

$$\begin{aligned}
\operatorname{Var}(H_0) &= t_0 + O(1), & \operatorname{Var}(Z_j) &= L + O(1),\\
\|\operatorname{Cov}(G_i, G_j)\|_{\mathcal{S}_2} &\le C\vartheta^{|i-j|} \quad (i \ne j),\\
\|\operatorname{Cov}(G_i, Z_j)\| &\le C\begin{cases}\vartheta^{i-j}, & j \le i,\\ \vartheta^{j-i-1}, & j \ge i+1,\end{cases}\\
|\operatorname{Cov}(Z_i, Z_j)| &\le C\vartheta^{|i-j|-1} \quad (i \ne j),\\
\|\operatorname{Cov}(G_i, H_0)\| &\le C\vartheta^{i} + O_K(r_i),\\
|\operatorname{Cov}(Z_j, H_0)| &\le C\vartheta^{j-1} + O_K(r_{j-1}).
\end{aligned} \tag{*}$$

The norms involving one scalar are norms of the corresponding linear functional on the normalized shape space.

For the shape–shape term, use the separated zero-mass logarithmic kernel estimate in [14, Lemma A.1, equation (35)]. Applied to the radius-8 pin and the fixed rescaled geometry, its cancellation estimate bounds every fixed-order derivative of the cross kernel by Cr_j/r_i for $i < j$. The Fourier calculation above therefore gives the Hilbert–Schmidt bound required here. For G_i, Z_j, both potential differences are smooth on the observed shape annulus. There are only two unsuppressed cases, $j = i$ and $j = i+1$. In every other case cancellation of the smaller increment or smaller shape gives the powers displayed in (*). For Z_i, Z_j with $|i-j| \ge 2$, the smaller increment lies strictly inside the smaller of the two spheres forming the larger increment; its zero total mass gives the same decay. For adjacent increments the leading $\min(t_i, t_j)$ terms in the four average covariances cancel, leaving $O(1)$. The same average-covariance calculation gives $\operatorname{Var}(Z_j) = L + O(1)$. Finally, subtracting the fixed exterior pin contributes only a smooth potential. Its oscillation against a scale-r zero-mass test is $O_K(r)$, which proves the last two estimates. Increasing the starting scale makes all these exterior errors uniformly negligible.

Let $\mathcal{R}$ be the complete covariance operator relative to the block-diagonal covariance of the product experiment. Summing the rows in (*) after the scalar normalizations gives

$$\|\mathcal{R} - I\| \le CL^{-1/2} + CN^{-1/2}, \qquad \|\mathcal{R} - I\|_{\mathcal{S}_2}^2 \le CM/L + C/N.$$

Choose L and then N so the first bound is less than $1/2$. In particular there is no hidden linear relation between the shapes and coarse history. If $1+\delta_k$ are finite-projection covariance eigenvalues, their Gaussian density has second moment

$$\prod_k (1-\delta_k^2)^{-1/2} \le \exp(C\sum_k \delta_k^2).$$

The uniformly bounded density martingale proves (24.1) in the infinite-dimensional limit.

For score-augmented shapes, every cross covariance involving a score and another shape, another score, or a coarse average is zero: it is the induced shift evaluated on a region or sphere where that shift vanishes. Independent white-noise times below all recipe bands in the given shape block give uniform coercivity of the shape-test covariance inside each augmented block, as in Lemma 23.6. Thus (*) and the same operator proof hold with a larger fixed constant. □

Lemma 24.2 (Uniformly bounded alignment shifts)**.** *Let ν be the unit law of one shape block, optionally augmented by the preceding finite score catalogue. Let $K_i(g,\cdot)$ be arbitrary local observation kernels and let E_i be events satisfying $(\nu\otimes K_i)(E_i^c)\le\epsilon$. For every $A<\infty$ there is a finite C_A such that, for every $|c|\le A$ and every i,*

$$(\nu^{(c)}\otimes K_i)(E_i^c) \le C_A\epsilon^{1/2}.$$

Here $\nu^{(c)}$ adds the constant c to the field on S and leaves every retained score unchanged. A fixed equivalent unit reference law may be used instead, with one modulus of absolute continuity uniform over i and the kernels.

Proof. Choose a smooth cutoff equal to one on a buffered neighbourhood of $\overline{S}$ and zero near the radius-8 pin and outside a slightly larger annulus. Use the localized white-noise shift construction at a time band strictly below all recipe-score bands. Convolution preserves the plateau exactly, so its induced field shift equals c on S, vanishes at the pin, and its inner product with every recipe direction is zero. Its squared white-noise energy is at most Cc^2. The Gaussian Radon–Nikodym density consequently has second moment at most e^{Cc^2}. Cauchy–Schwarz proves the claim. Adjoining the same kernel leaves the density unchanged. For an equivalent reference law, truncate its fixed density before applying this estimate. The truncation error is uniform over all kernels and indices. □

Lemma 24.3 (Exact physical index identities before regular variation)**.** *Let $I_n\to\infty$ and let $B_n>0$ be deterministic cost units. For an actual field input g, let $\mathcal{Q}_n(g,\cdot)$ be the unnormalized Poisson passage-record kernel at intensity $I_n\mu_g$. Set $v_i=b^{-i}$, $c_i=\gamma^{-1}\log v_i$, and define the indexed experiments by*

$$K_{n,i}(g,\cdot) = (S_{(B_n v_i^\alpha)^{-1}})_*\mathcal{Q}_n(g+c_i,\cdot). \tag{QI1}$$

Suppose their unit local joint laws have been extracted with all fixed indices, fixed affine charts and background coordinates, using the fixed-density local-kernel closure. Then their limiting kernels satisfy

$$K_i(g,\cdot) = (S_{v_i^{-\alpha}})_* K_0(g+c_i,\cdot), \tag{QI2}$$

and the actual physical restriction on $z+rS$, expressed in unit coordinates and divided by $B_n v_i^\alpha$ before extraction, has limiting kernel

$$K_i\big(h(z+r\,\cdot)+Q\log r - c_i,\cdot\big). \tag{QI3}$$

These statements hold jointly with the complete coarse history and the fixed white-noise background. They neither assert regular variation of B_n nor a Weyl rule for the candidate metric. Kernel evaluations in these formulas mean their common local versions on admissible observation patches, as almost-sure conditional-kernel identities under the displayed Gaussian inputs. They are not pointwise identities for a global kernel evaluated on a new singularly pinned field law.

Proof. For a similarity $\psi(x) = z + rx$, the canonical GMC convention gives the exact measure identity

$$\mu_h(\psi(A)) = r^{d+\gamma^2/2}\mu_{h\circ\psi}(A) = \mu_{h\circ\psi+Q\log r}(A). \tag{QI4}$$

Pulling the Poisson cloud back by ψ therefore gives the Poisson construction with that last field input at the same intensity I_n. Euclidean Voronoi diagrams and their Delaunay adjacency are similarity covariant, and all unit-hop and distinct-cell cover costs are unchanged by this pullback. Guards and testing ports, when present, are pulled back with the chart. Moreover $\mu_{g+c_i} = v_i\mu_g$, so (QI1) is exactly the intensity-v_iI_n experiment divided by the *prescribed* unit $B_nv_i^\alpha$. This unit is not declared to equal $b(v_iI_n)$. Substituting $g = h\circ\psi + Q\log r - c_i$ in (QI1) proves the prelimit version of (QI3), and (QI2) is the same constant-shift identity relative to the base unit B_n.

For the passage to the limit, fix one chart and one index first. The physical input $Y = h\circ\psi + Q\log r - c_i$ on S has one fixed Gaussian law, independent of n. On the admissible observation window used here, that law is equivalent to the unit shape law ν. Its density is a fixed integrable function of Y. Approximate that density in L^1 by bounded continuous cylinder functions, as in Proposition 23.5. This transfers the extracted unit-law kernel to the actual law of Y. To retain any bounded background test Φ, use instead the fixed integrable weight

$$\frac{d\mathcal{L}(Y)}{d\nu}(Y)\mathbb{E}[\Phi \mid Y].$$

The same approximation proves the identity jointly with every coarse average, the white noise and the actual local reference readouts. The argument for the deterministic shift c_i proves (QI2). A countable diagonal gives all fixed charts and indices together.

In the annular application the input Y is merely rewritten as

$$Y = G_j + c, \qquad c = H_j - Qt_j + \gamma^{-1}i\log b. \tag{QI5}$$

No pointwise convergence of $K_{n,i}(G_j + c)$, and no continuity of a limiting kernel in c, is being asserted. Its joint conditional law has already been identified at the actual Gaussian input Y with the whole history retained. The subsequent Gaussian comparison replaces the joint shape/history law by its product law while appending these same Borel local kernels. Under that product law, conditioning on the history makes c a fixed constant; the bounded alignment lemma then controls its marginal event probability. A history-measurable choice among finitely many indices is valid because the identities hold simultaneously for those indices conditional on the full field.

For every fixed label window N, all potential centres, radii, labels and shifts belong to one finite part of the countable library. The discrete limit $n \to \infty$ is taken first, and only afterwards is $N \to \infty$ used on that retained limiting array. The proof therefore makes no uncontrolled moving-chart extraction. The cone estimates supply tightness in the units $B_nv_i^\alpha$; they are not used as a multiplicativity identity for the normalizer. □

Proposition 24.4 (All-point annuli at deterministic quantum labels)**.** *Let $d \geq 2$, $0 < \gamma < \sqrt{2d}$, $Q = d/\gamma + \gamma/2$, and $b > 1$. Suppose the local passage records have the conditional product kernels, fixed-law compatibility, and actual coordinate/index identities described in the restricted interface:*

at volume $v_i = b^{-i}$ the physical record on $z + rS$, divided by its indexed cost unit, has kernel K_i at input

$$h(z + r\,\cdot) + Q \log r + \gamma^{-1} i \log b.$$

All fixed indices needed in one extraction are retained together. The records may include the finite localized-score catalogue.

For every $\theta_0 > 0$ there are $0 < a < B$, $A < \infty$, $\beta > 0$, $\epsilon_* > 0$, and $c, C > 0$ with the following property. Let K be a fixed physical compact set. For every sufficiently large N, choose a deterministic set $I_N \subset \{N, \dots, 2N\}$ with $|I_N| \geq \theta_0 N$. For $i \in I_N$, let E_i be a unit local event of probability at least $1 - \epsilon_*$. Except with probability Ce^{-cN}, every $x \in K$ belongs to the central holes of at least βN pairwise disjoint buffered annuli passing their assigned tests. Their radii lie on one fixed geometric sequence between e^{-BN} and e^{-aN}, their centres lie on deterministic meshes of spacing at most $c_* r$, their labels are distinct for each x, and

$$\left|\gamma(Q \log r + H(z, r)) + i \log b\right| \leq A.$$

The constants may be chosen independently of the physical compact after increasing its starting N. The conclusion also holds for $v_i = Mb^{-i}$ with any fixed $M > 0$ after changing that starting value.

Proof. Choose $\eta > 0$ with $Q > \sqrt{2d} + \eta$, and choose a, B such that

$$\gamma(Q + \sqrt{2d} + \eta)a < \tfrac{1}{2} \log b, \qquad \gamma(Q - \sqrt{2d} - \eta)B > 3 \log b.$$

Take logarithmic scales on the fixed progression of spacing L from Lemma 24.1, restricted to $[aN, BN]$. At each scale use a mesh of spacing at most $c_* r$. Each point of the finest mesh follows its nearest centres on the coarser meshes, giving a deterministic ancestry. Decrease c_* and increase L so its buffered annuli are disjoint and every finer normalization sphere lies in every preceding central hole.

At an endpoint scale t, there are $O_K(e^{dt})$ possible centres, and $\operatorname{Var} H = t + O_K(1)$. A Gaussian union bound therefore gives, simultaneously at both endpoint scales and all their centres,

$$|H| \leq (\sqrt{2d} + \eta)t$$

except with probability $C_K e^{-cN}$. Notice that the endpoint mesh is counted at its own scale. On this event $\Lambda_j = \gamma(Qt_j - H_j)$ starts below $N \log b$ and ends above $2N \log b$ for every ancestry.

Assign a target $i \log b$ to its first upward crossing by Λ_j. No monotonicity is required. A jump of size $s > A$, where $A \geq 2 \log b$, receives at most $2s/\log b$ targets. Hence the number assigned to such large jumps is at most

$$\frac{2}{A \log b} \sum_j |\Lambda_j - \Lambda_{j-1}|^2.$$

The covariance operator of the increment vector $(Z_j)_j$ is bounded by CL, and its trace by CLM, by (*). Thus

$$\mathbb{E} \exp\{cL^{-1} \textstyle\sum_j Z_j^2\} \leq e^{CM}.$$

Since $\Lambda_j - \Lambda_{j-1} = \gamma QL - \gamma Z_j$, choosing A sufficiently large makes the probability of losing more than $\theta_0 N/4$ labels at most e^{-qN} for any prescribed fixed q. This choice is made once, before the local failure-probability bound. Every remaining assignment has alignment error at most A, and one annulus receives at most $J = \lceil A/\log b \rceil + 2$ labels. Retain one target from each such annulus. At least νN trials remain, where $\nu > 0$ depends only on the fixed parameters.

Compare the joint shape/history law for this ancestry with the product experiment of Lemma 24.1. Conditional on the coarse variables in the product experiment, the selected labels and shifts are deterministic. The shapes are independent, and the actual coordinate identity gives exactly the input

$$G_j + H_j - Qt_j + \gamma^{-1} i \log b.$$

The constant shift has magnitude at most A/γ. Lemma 24.2 allows ϵ_* to be chosen so every selected conditional failure probability is at most any prescribed $\zeta > 0$. The product local kernels make the trial outcomes independent conditional on those coarse variables. Consequently the product probability of fewer than $\nu N/2$ successes, on a coarse history with at least νN retained trials, is at most

$$2^{M+1} \zeta^{\nu N/2}.$$

Restore the actual Gaussian law using Cauchy–Schwarz and (24.1). The actual probability is at most

$$e^{CM/(2L)+C}\, 2^{(M+1)/2} \zeta^{\nu N/4}.$$

Choose $q > dB + 3$ and then ζ small enough that this is at most e^{-qN}. There are only $O_K(e^{dBN})$ finest-mesh ancestries. The union bound for the large-jump and trial failures is therefore exponentially small. The already global endpoint failure has its own exponentially small bound and is not multiplied by the finest mesh count.

Every point of K is close enough to a finest-mesh point to lie in all its selected central holes, after one last fixed reduction of the mesh spacing. Set $\beta = \nu/2$. The argument compares each ancestry's actual law; it does not require product couplings to be consistent across overlapping ancestries. Replacing b^{-i} by Mb^{-i} shifts every target by the same constant $-\log M$, which the endpoint margins absorb. □

Remark 24.5 (Scope and exact role of the new comparison). The proposition replaces the planar use of harmonic extensions and independent Dirichlet fields. Its local-kernel hypothesis follows from the preceding fixed Gaussian density argument for the retained discrete records; its coordinate/index identity is an exact transported-experiment identity, not a Weyl rule for the candidate passage law. A fixed locally equivalent field change preserves the eventual coverage conclusion along a summable sequence of windows, although its exponential rate need not be preserved. This result provides the probability transfer needed to initialize finite comparison by high-probability annular tests. The deterministic comparison step and the later forcing argument remain separate.

Remark 24.6 (A fixed pin and coverage of the whole space). A single global pin cannot be placed outside every member of an increasing compact exhaustion. The whole-space use of the proposition is instead made through the admissible atlas of Lemma 21.1. Choose a compact $K \Subset W$ with a nonzero part of the smooth pin outside $\overline{W}$. For all sufficiently large N, the outermost radius-8 normalization spheres, every smaller history sphere, all the observed annuli and all their auxiliary-record windows lie inside W: their largest radius is $O(e^{-aN})$, while K has positive clearance in W. Thus no ancestry history outside W is read. The finitely many omitted initial windows play no role in arbitrarily fine coverage.

The field law on W is equivalent to that of a whole-space field whose smooth pin lies outside $\overline{W}$. Transfer the actual local arrays and their common local reference readouts through this one fixed density, as in Proposition 21.2. The event of eventual coverage along a summable sequence of N is measurable from the countable local field-and-kernel array on W, and has probability one under the exterior-pinned law. Equivalence transfers that event to the actual fixed-pin law. This transfer keeps the numerical comparison constants; it need not keep an exponential failure rate. A

countable family of such smaller compacts covers the whole space. At no step is an unidentified candidate metric multiplied by a random Weyl factor or evaluated at a random intensity.

For this atlas transfer, use the unaugmented local field-and-record coverage event. Retained raw white-noise recipe scores need not be functions of the local field, so a density of that field alone is not asserted to transfer arbitrary score events. Once coverage has been transferred, impose likelihood clips under the original white noise by Lemma 24.7. Its separate all-ancestry deletion estimate does not condition the scores on the coverage event.

Retaining the full white noise and the reference readouts in the joint limiting array does not mean that the Gaussian comparison above conditions on that full background. The probability proof uses the unconditioned finite shape/history Gaussian vector. Given that vector, the local observation kernels depend on the displayed actual local field inputs; extra background coordinates are integrated out. Full-background retention is used separately in the exact completion proposition, where one conditions on the white noise to identify the passage kernels, not to assert independence of the Gaussian shape variables.

Lemma 24.7 (Score clipping after the catalogue has been chosen)**.** *Apply Proposition 24.4 without adjoining recipe scores, and let $\beta > 0$ be its success density. After its geometric scale sequence and any finite deterministic recipe catalogue have been fixed, suppose every recipe direction is supported in a fixed buffered annular space–time region whose physical copies at distinct scales of one ancestry are disjoint. Their induced shifts vanish on the normalization spheres. For any $q > 0$ one can choose a finite clipping level L_{clip} such that, except with probability $C_K e^{-qN}$, every ancestry has at most $\beta N/2$ annuli on which any recipe score exceeds L_{clip} in absolute value.*

Consequently the original all-point coverage retains at least $\beta N/2$ successes per point after requiring all likelihood clips. No independence between catalogue success and score clipping is needed. The all-point event's original failure threshold can be chosen before the recipe catalogue is constructed.

Proof. At one annulus there are finitely many centred Gaussian scores, with finite variances equal to the recipe energies. Translation and dilation preserve their joint law. A Gaussian union bound makes the probability p of any clipping failure as small as desired by increasing L_{clip}. For a fixed ancestry, the entire score vectors at its different scales are independent, because their white-noise supports are disjoint. This uses only their common buffered annular support, not a common lower time cutoff across the scales or a covariance comparison with the field shapes.

There are at most $C_0 N$ scales. The probability that at least $\beta N/2$ of them fail clipping is at most

$$2^{C_0 N} p^{\beta N/2}.$$

There are at most $C_K e^{dBN}$ finest-mesh ancestries. Choose p so small that $C_0 \log 2 + (\beta/2) \log p + dB < -q$ and take their union bound. Each successful annulus in the original construction uses one scale of the relevant ancestry, so removing all clipped scales deletes at most $\beta N/2$ successes. This deterministic deletion bound is valid regardless of dependence between a score and the field-and-kernel event at its own annulus. □

Remark 24.8 (Why the separate clipping order matters)*.* The score-augmented Gaussian comparison is valid for any fixed finite catalogue, but its constants can depend on the smallest recipe time band. Constructing a catalogue after choosing that comparison's failure threshold could therefore introduce a circular choice. The preceding lemma avoids it: choose the unaugmented all-point threshold, construct the field-and-auxiliary-record catalogue to that threshold, and only then choose the score clipping level. Its binomial estimate is uniform in the scale count and does not alter the already fixed all-point geometry.

25 Annular comparison for passage records

This supplement isolates the deterministic step in the physical-label interface (L). It assumes that the good annular covers have already been obtained. Their balls and weights may depend on the realized field and passage array. The proof is pointwise on the event carrying the simultaneous countable record interface. It uses neither independence of the balls nor a parametrization of a limiting passage record range.

Let D be a continuous length metric inducing the Euclidean topology. The candidate object consists of compact-range passage records, with nonnegative costs, actual embedded-walk attainability in strict open buffers, bounded-cost subwalk capture, finite ordered marks, concatenation, and a local endpoint-joining modulus. The actual one-edge costs and spatial mesh tend to zero. At this stage a zero-cost passage record may have a nontrivial connected range, and that range need not be path connected.

25.1 Finite chains and the order of the cuts

Lemma 25.1 (A chain made only of charged balls)**.** *Let $B_i = \overline{B}(z_i, R_i)$ be a finite family of Euclidean balls of diameter at most 2ϵ, in dimension $d \geq 2$. Their open central holes $H_i \subset B(z_i, \theta R_i)$, with fixed $\theta < 1$, cover the range of a continuous path P. Assign each ball a weight $w_i > 0$. Suppose every pair of points of ∂B_i has a join of cost at most Aw_i, with arbitrary positive allowance, in a prescribed open annular neighborhood contained in $B(z_i, \Lambda R_i)$.*

Choose a covering hole at the current path point, follow the path to its first subsequent hit of that ball's boundary, and repeat. A ball with no subsequent boundary hit is the final incomplete ball. If the completed traversals have nonnegative charges q_k with $q_k \geq aw_{i_k}$, then there is a joined connector between points within 4ϵ of the original endpoints, of cost at most

$$A \sum_k w_{i_k} \leq \frac{A}{a} \sum_k q_k,$$

up to an arbitrary positive allowance. Its range is contained in a $C_\Lambda \epsilon$-neighborhood of the original path range. No comparison between the weights of nested balls is required.

Proof. Completed traversal intervals are disjoint in path order. Every one starts in its selected hole and ends on its selected boundary. Consecutive completed balls intersect at the preceding stopping point. For a continuous path and a finite family with central holes, the greedy procedure terminates: a completed traversal moves a fixed positive Euclidean distance, and uniform continuity of P prevents infinitely many such moves in its compact parameter interval. When a total cost is available, the lower charges also give the useful bound on the number of completed traversals.

Consider only balls that appeared in completed traversals. Replace each by a maximal member of this finite completed list containing it. Consecutive replacements still intersect. Retain a simple chain in the resulting intersection graph from a maximal ball containing the initial completed ball to one containing the final completed ball. Distinct maximal balls do not contain each other. Two intersecting Euclidean balls with neither containing the other have intersecting boundaries when $d \geq 2$. Choose a common boundary point at every link of the simple chain and use the specified annular joins within the successive balls.

Each retained ball is an actually completed ball, and no retained ball is used twice. Its weight was therefore included at least once in the original sum over completed traversals. This proves the cost bound without relating different balls' radii or weights. It would be unjustified to insert an untraversed maximal ball from the original cover; the construction does not do that.

The first retained ball contains the initial path point. The last retained ball contains the last completed stopping point, whose distance from the final endpoint is at most 2ϵ if there is a final incomplete ball. Choosing initial and final points on the retained boundary spheres therefore gives endpoint errors at most 4ϵ. If no traversal was completed, the entire path lies in its final incomplete ball and a constant connector suffices.

Every retained ball itself was traversed, so its central hole met the original path. Its prescribed annular join lies within $C_\Lambda \epsilon$ of that path. This proves confinement of the whole connector. In particular, if the original range is compactly contained in an open set, all these joins lie in that open set for sufficiently small ϵ. □

Lemma 25.2 (A limiting shell charge is uniform before counting cuts)**.** *Fix one finite family of balls and positive weights as above in an almost surely convergent representation of the actual record arrays. Suppose every limiting passage record crossing the prescribed closed interior shell of ball i, with range in that shell, has cost at least $aw_i > 0$. The shell is compactly contained in its observation chart. Then for every $0 < a' < a$, for all sufficiently large discrete indices, every actual completed hole-to-boundary traversal has normalized cost at least $a'w_i$.*

Consequently, actual walks with a common total cost bound have a uniformly bounded number of completed traversals in this fixed cover. This assertion does not assume any positive global lower comparison for the candidate object.

Proof. If the claim failed, choose a violating traversal along an infinite subsequence and pass to one fixed ball. Cut its actual walk from the last visit to the inner testing sphere before the first visit to the outer testing sphere. The continuous crossing range lies in the closed intervening shell. Replace cut points on graph edges by adjacent vertices. This changes spatial positions by the vanishing mesh and the cost by at most two vanishing one-edge costs. Thus the resulting bounded-cost subwalks lie in one fixed compact observation chart, have endpoints tending to the two testing faces, and have limiting cost at most $a'w_i < aw_i$.

Capture this one subwalk. Its limiting range is contained in the closed shell and its endpoints are on the required faces, so it contradicts the assumed shell charge. The same argument applies to a rectangular slab inside the traversal, using its last-inner and first-outer face cuts. Finiteness of the family makes the eventual claim simultaneous for all its balls.

The traversal intervals are disjoint before any deletion of balls. Their costs are measured on the actual walk, so their sum is at most its total normalized cost. Dividing by $a' \min_i w_i > 0$ bounds their number. Only after this bound is proved need one capture the complete finite list of ordered cuts; all its edge errors then vanish. There is no stopping-time assumption and no circular appeal to a limiting rectifiable path. □

25.2 The two comparisons

Theorem 25.3 (Two-sided annular transfer for full passage record objects)**.** *Suppose every compact set has arbitrarily fine finite covers by open central holes of balls from the countable inspected library. Every corresponding annulus lies in a fixed multiple of its ball, and its maximum radius tends to zero. On each cover, assign weights $w_B > 0$.*

For the upper comparison assume that every reference traversal of the tested shell has D-length at least mw_B, and every pair of boundary points has a realizable candidate join, confined to the prescribed annular neighborhood, of cost at most Mw_B. Then, for every connected open U,

$$F_U(x, y) \le \frac{M}{m} D(x, y; U).$$

More precisely every reference-rectifiable path compactly confined in U admits candidate follows in any prescribed open neighborhood of its range with this cost factor and arbitrary positive allowance.

For the lower comparison instead assume that every candidate shell crossing passage record has cost at least mw_B, and every pair of boundary points has a reference join inside the prescribed annular neighborhood of length at most Mw_B. Then every finite-cost passage record with endpoints x, y and compact range in U has cost at least

$$\frac{m}{M} D(x, y; U).$$

The lower conclusion is valid when the old optimal lower comparison constant is zero and before candidate noncollapse is known. Fixed changes of the buffered testing geometry merely alter the deterministic geometric constant in these comparisons.

Proof. For the upper statement fix a finite-length reference path P compactly in U. Use a fine cover of a deterministic compact containing its range and the greedy stopping rule. The reference lengths of completed traversal intervals are disjointly charged, so

$$m \sum_B w_B \le \mathrm{len}_D(P),$$

where the sum includes traversal multiplicities. Apply Lemma 25.1 to the candidate boundary joins. For this fixed finite cover, choose their finitely many allowances with arbitrarily small total and concatenate their passage records. All ranges lie in a vanishing Euclidean neighborhood of P. The preexisting candidate endpoint-joining modulus corrects the two endpoint errors inside U, with cost tending to zero as the cover is refined. This gives the asserted upper follow and, after taking the infimum over P, the intrinsic upper comparison. It does not require a candidate metric lower bound.

For the lower statement fix a passage record of total upper cost c and compact range $K \Subset U$. First fix a sufficiently fine finite cover by open holes. Their union covers a whole neighborhood of K, so every sufficiently late actual realizing walk is covered by those holes. Choose realizing walks whose costs tend from above to at most c, endpoints tend to x, y, and ranges tend to K. Apply Lemma 25.2 with $m' < m$. For all sufficiently late walks,

$$m' \sum_B w_B \le c + o(1).$$

Now apply Lemma 25.1 to the reference boundary joins. This produces a reference connector in a small neighborhood of K of length at most $M(c + o(1))/m'$, between points within the prescribed small Euclidean endpoint error of x, y. Reference local continuity and the length property correct those endpoint errors inside U. First take the discrete limit with this cover fixed. Then refine the cover, and finally let $m' \uparrow m$. The result is $D(x, y; U) \le Mc/m$.

No actual graph cut count growing with the discrete resolution has been passed through a finite-mark compactness argument. At each fixed cover the positive shell charge bounded the count first. Also no comparison on the boundary of an observation chart was used: the cut shell is strictly inside that chart, and the final reference joins lie in the given open domain.

The construction is simultaneous on the event carrying the countable library's capture and joining assertions. A finite cover may consequently be selected from that library after the limiting field or passage record range is known. This selection is only part of a deterministic path argument, not a conditioning operation. □

Remark 25.4 (Full ranges, ordered marks, and zero-cost passage records). If a follow must retain prescribed points or subranges of a reference path, apply the upper construction to its finitely many ordered subpaths before concatenating. Refine the marked partition and then the covering radius. The marked points become dense in the old range and all joining ranges remain close to it, so capture recovers the full range and its ordered subpiece cost bounds. Removing loops from one finite ball chain is therefore not a claim that the full reference path's traversal order is automatically preserved by that chain.

For a candidate passage record of zero cost, the fixed-cover argument eventually permits no completed traversal. Its realizing walks then lie in a single ball of that cover. Arbitrarily fine covers force the retained range to be a singleton. More generally, once the positive lower comparison is known, every marked subrange is controlled by its own cost, which is exactly the premise needed for the subsequent intrinsic path-calculus construction.

25.3 The high and low amplification calculations

Lemma 25.5 (Savings and extra charges, including a zero old lower factor). *Suppose $0 \le c < C < \infty$ are valid old lower and upper comparison constants. Fix a finite good cover with the preceding confinements and reference surrounding joins of length at most $A_0 w_B$.*

For the high calculation, let P be a reference geodesic minimizing in an open domain containing all the selected annular joins. Suppose each completed traversal has a disjoint reference geodesic subarc of length at least $a_0 w_B$, admitting a candidate follow at factor at most $C - \delta$. Then, after endpoint corrections and refinement of covers, the candidate upper factor improves to every constant greater than $C - \delta a_0/A_0$.

For the low calculation, suppose one disjoint marked slab piece in each completed traversal has candidate cost at least $c_1 D(x_B, y_B; R_B^+)$, where $c_1 > c$ and the denominator is at least $a_0 w_B$. The remaining pieces retain the old lower comparison. Then the lower factor improves to every constant smaller than $c + (c_1 - c)a_0/A_0$. This holds also for $c = 0$.

Proof. For the high calculation, follow the charged subarcs at the better factor and all remaining pieces at factor C. They are disjoint in the original geodesic's time, giving total candidate cost at most

$$C \operatorname{len}_D(P) - \delta a_0 \sum_B w_B + o(1).$$

The reference version of the finite ball chain connects near the endpoints inside the minimizing domain. Since P is a geodesic in that domain,

$$\operatorname{len}_D(P) \le A_0 \sum_B w_B + o(1).$$

This gives the claimed improvement. The minimizing hypothesis is essential for this last inequality; it is not true for an arbitrary rectifiable path. To extend the improved comparison to general paths, first subdivide them into short pieces and replace each by a reference minimizing connection in a small interior neighborhood. Alternatively use internal geodesics compactly confined in the requested domain. All candidate follows keep strict domain margins.

For the low calculation, use actual realizing walks. Each marked slab piece has the strictly positive charge $c_1 a_0 w_B$, so the prelimit-charge lemma bounds the number of traversals even when $c = 0$. Capture a finite ordered decomposition into these pieces and the remaining pieces. On a slab piece write

$$c_1 D(x_B, y_B; R_B^+) \ge c D(x_B, y_B) + (c_1 - c) a_0 w_B.$$

The remaining pieces have cost at least c times their reference endpoint distance. Summing in order and using the reference triangle inequality gives

$$c_{\text{total}} \geq cD(x,y) + (c_1 - c)a_0 \sum_B w_B - o(1).$$

If an intrinsic conclusion is required and $c > 0$, use the old intrinsic comparisons in the common open domain at this step; all slab boundary layers and remaining subpath ranges lie there. When $c = 0$, only nonnegativity is used for the remaining pieces. The reference finite ball chain gives $D(x,y) \leq A_0 \sum_B w_B + o(1)$, or the same bound for the internal distance in that domain. This proves the claimed strict improvement. One may first replace c_1 by a number strictly between c and c_1 to leave capture margins. □

25.4 Source and output arclength charges

Lemma 25.6 (Source coverage and output multiplicity)**.** *Let P be a reference geodesic and let P_{mid} be a compact strict interior subarc. Suppose every point of that subarc lies in at least βN good central holes. The corresponding full balls are small enough to miss both endpoints of P, and any two points on each surrounding boundary sphere have a reference join of length at most Aw_B in the geodesic's minimizing domain. Then*

$$\sum_{B:\, P_{\text{mid}} \cap H_B \neq \varnothing} w_B \geq \frac{\beta N}{A} \operatorname{len}_D(P_{\text{mid}}).$$

For an output state use its one fixed-anchor reference geodesic P^{out}, with length at most L_*. Suppose output test B_β charges at least κw_β to a subarc of that geodesic contained in an enlarged ball $\widehat{B}_\beta$. At each Euclidean radius the enlarged balls have bounded overlap, each center/radius has a bounded number of possible labels and recipes, and there are $O(N)$ radii. Then, pointwise in the output,

$$\sum_\beta w_\beta \mathbf{1}_{B_\beta} \leq KNL_*.$$

The latter inequality also permits output-measurable random weights.

Proof. For a source ball hit by the middle arc, take the first and last visits of the full geodesic to its surrounding boundary sphere. The endpoints of the full geodesic are outside the ball, so these visits exist. The intervening subpath is minimizing in the domain where the stipulated boundary join is allowed, and therefore has length at most Aw_B. It contains every visit to that ball's central hole. This proves

$$\operatorname{len}_D(P_{\text{mid}} \cap H_B) \leq Aw_B,$$

where the left side means occupation in the arclength parameter, including repeated visits. Integrate the pointwise multiplicity $\sum_B \mathbf{1}_{H_B}(P(t)) \geq \beta N$ over the middle arc and sum the preceding bounds. No single-visit assumption is made.

For the output, the retained crossing subarc gives

$$\kappa w_\beta \mathbf{1}_{B_\beta} \leq \int_{P^{\text{out}}} \mathbf{1}_{\widehat{B}_\beta}\, d\ell_D.$$

At a fixed radius, summing the right side uses the bounded overlap and the bounded number of label/recipe copies. It is at most a fixed multiple of the length of this same output geodesic.

Summing over the $O(N)$ radii proves the pointwise estimate. Different outputs may have different geodesics; within a single output state, the geodesic in every summand is the same fixed-anchor readout. If the weights are random, the estimate remains pointwise and they must subsequently be integrated inside the output submeasures.

In the physical-label application, bounded alignment bounds the number of labels at a fixed center and radius, and the recipe list is finite. The further-shell passage record supplies the required subarc by last-inner/first-outer cutting. It is read in the output metric; a source-only lower shell bound would not justify this charge after modification. □

26 Spatial amplification of a confined shortcut in every dimension

The exact-shift fresh-sample mechanism follows [1, Proposition 8.26, pp. 89–91]; its required Gaussian separation is verified here.

The forcing catalogue needs many independent spatial opportunities for a low-ratio passage record. Repeated auxiliary samples at one fixed field would not amplify an event that fails because of that field. The following construction uses independent fine white-noise fields on separated trial windows. It is valid in odd as well as even dimension and does not use a domain Markov property.

Let the whole-space logarithmic field be realized by the compact kernel k and white noise W used in the preceding localized-shift construction. Fix a bounded unit dependency window D, including every denominator boundary layer and preservation patch needed for a trial. Its reference field is pinned by a smooth average outside a larger window. Let $K_i(g,\cdot)$ be any family of local observation kernels, and suppose a local trial event E_i satisfies

$$\mathbb{P}_{\mathrm{ref}}\big((h,A_i)\in E_i\big)\ge\beta>0,\qquad A_i\mid h\sim K_i(h,\cdot),\tag{SA1}$$

uniformly over the favourable indices i. A finite union of endpoint bins or geometric types may be used as E_i; no individual bin needs a uniform positive probability.

Proposition 26.1 (Independent fine-field trial amplification)**.** *Fix an open trial region in a unit shell, separated from its central hole and outer boundary, and fix $\theta>0$. There are finitely many deterministic disjoint affine trial charts $\psi_j(u)=z_j+\rho u$, and, for every fixed $c\in\mathbb{R}$, a finite deterministic set of exact shifts $\mathcal{S}_c\subset\mathbb{R}$, such that the following holds uniformly over the favourable indices. For every pair (j,s) draw an auxiliary local record from*

$$K_i\big(h\circ\psi_j+Q\log\rho+c+s,\cdot\big),\qquad s\in\mathcal{S}_c,\tag{SA2}$$

using independent auxiliary seeds for different pairs, conditional on the field. With probability at least $1-\theta$, at least one of these exact-shift records and its actual local field satisfy E_i. The successful-trial event is measurable from the field restricted to the outer observation shell and the finite auxiliary record array. The constants and grid do not depend on i.

In the forcing application $c=\gamma^{-1}\log M$, and (SA2) is the exact transported kernel obtained from the outer experiment at volume Mv_i after dividing by the finer quantum cost unit. This interpretation uses the index/coordinate identity; it assumes no Weyl scaling for the candidate passage costs.

Proof. Fix a large deterministic ℓ and split white noise at time $T=\ell\rho$. The kernel has compact support, so the fine fields on two trial windows use disjoint noise whenever their enlarged dependency regions, of radius $O(\ell\rho)$, are disjoint. A finite number of such windows can be packed in the prescribed region by taking ρ small enough.

For clarity about the global normalization, place its smooth pin outside the outer shell and choose T smaller than its separation from that shell. Condition on all noise outside the disjoint local fine cones. This includes the low-time contribution to the exterior pin. In each trial chart the remaining field has the exact form

$$h \circ \psi_j = X_j^{(\ell)} + H_j + \Phi_j, \qquad \Phi_j(0) = 0, \tag{SA3}$$

where the $X_j^{(\ell)}$ are independent, identically distributed finite-range logarithmic fields with time cutoff ℓ, independent of the conditioned background. The scalar H_j and the smooth profile Φ_j are background measurable. On the dependency window the profile is the oscillation of the noise with times at least T.

For every fixed integer k, its derivatives satisfy the uniform estimate

$$\mathbb{E}\|\Phi_j\|_{H^k(D')}^2 \le C_{k,D',\ell}, \tag{SA4}$$

on a slightly larger fixed window D'. This bound is independent of ρ, z_j, and the number of trials. Indeed, for every derivative order $m \ge 1$, its pointwise variance is bounded by

$$C_m \rho^{2m} \int_{\ell\rho}^{\infty} t^{-2m-1}\, dt \le C_m \ell^{-2m},$$

and the zero-order difference is controlled by the first derivative. The same estimate after a fixed number of additional derivatives controls the displayed Sobolev norm.

The law of $X^{(\ell)}|_D$ is equivalent to the fixed reference law in (SA1). Both local covariance norms are equivalent to $H^{-d/2}$; their covariance difference is smooth, so the normalized difference is Hilbert–Schmidt, and the local covariance operators are strictly positive. The finite-projection density argument used earlier proves equivalence. Truncating this one fixed density shows that (SA1) has probability at least a number $\beta_2 > 0$ under $X^{(\ell)}$, uniformly over i and all the observation kernels.

Choose $k > d/2$ and then R so large that (SA4) gives

$$\mathbb{P}[\|\Phi_j\|_{H^k(D')} > R] < \theta/100.$$

A fixed cutoff extends $\Phi_j + a$ on D, for $|a| \le 1$, to a whole-space Cameron–Martin function for $X^{(\ell)}$ with squared energy at most $C(R+1)^2$. To see this directly, the spectral covariance of the cutoff field is comparable to $(1+|\zeta|^2)^{-d/2}$, with constants depending only on ℓ. Thus its Cameron–Martin energy is bounded by a constant times the $H^{d/2}$ norm of the cutoff extension, which is controlled by its H^k norm. The inverse Gaussian likelihood inequality yields the conditional lower bound

$$\mathbb{P}\big((X_j^{(\ell)} + \Phi_j + a, A_i) \in E_i \mid \text{background}\big) \ge \beta_* := \beta_2^2 e^{-C(R+1)^2} > 0 \tag{SA5}$$

on every acceptable profile. Appending the same observation kernel does not change this likelihood calculation.

Choose the number J of trials so that $(1-\beta_*)^{J/2} < \theta/10$, and then choose and pack the J deterministic trial charts with their enlarged fine cones disjoint. The expected fraction of unacceptable profiles is below $\theta/100$, so Markov's inequality makes the probability that more than half are unacceptable below $\theta/50$. This step requires no independence among the coarse profiles.

The trial scale and centres are now fixed. The finite scalar vector $(H_j + Q\log\rho + c)_{j=1}^J$ is tight. Truncate it to a deterministic compact interval with failure probability below $\theta/10$, and

choose a finite exact-shift grid $\mathcal{S}_c$ whose spacing is at most two and which covers the negatives of that interval. For each trial choose, for the purpose of the probability estimate only, the nearest grid value. Conditional on the background this choice is deterministic, and its residual constant has magnitude at most one. Therefore (SA5) applies. The fine fields are independent and their auxiliary seeds are independent, so the probability that all acceptable chosen-grid trials fail is at most $(1-\beta_*)^{J/2}$. Summing the three error bounds gives the claim, after reducing the initial allocations if necessary.

The final success event simply asks whether any member of the finite array in (SA2) satisfies its local test. It does not expose the coarse/fine decomposition, the profile truncations, or the selected nearest grid values. Those background quantities were used only to prove its probability bound. Every test in that event is therefore measurable from the actual field in its dependency window and its auxiliary local record, which proves the claimed locality. □

Remark 26.2 (Order of choices in the forcing catalogue). Prepare the low event, denominator truncations, endpoint bins, and safe terminal caps first. Apply the proposition to choose the profile bound, number and placement of trials. Choose the fixed quantum separation M to make the finer reference competitor cheaper than the outer crossing lower bound. Only then truncate the finite compensation vector and choose the finite exact-shift grid. Draw the complete array before choosing a successful trial or a recipe. The denominator may be read on a larger source boundary layer, while only the passage record restricted to its preservation patch is retained in the target completion. Thus the proof does not identify a source denominator with the metric of a recipe which is constant only on the smaller preservation patch.

Remark 26.3 (Adding the remaining local catalogue tests). Reference crossing lower bounds, internal connection upper bounds, and local moduli can be intersected with this success event at an arbitrarily small fixed loss of probability. After the finite geometry and all deterministic localized white-noise recipes have been fixed, clip their finitely many Gaussian scores at a common deterministic threshold. Their variances are their dilation-invariant white-noise energies, so a Gaussian union bound makes this further failure probability arbitrarily small. Apply the all-point annulus proposition first to the unaugmented field-and-array event, with its failure threshold chosen before the recipes. Then use the separate binomial estimate of Lemma 24.7: the score vectors on disjoint ancestry annuli are independent, so one clipping level deletes at most half of the successful scales simultaneously over the finest-mesh ancestries. No independence between a local success event and its own scores is needed, and no recipe-dependent augmented threshold is fed back into the earlier choice of catalogue. The explicit wall, tube, and entrance set geometry and the deterministic shortcut implication must still be checked separately.

27 Corridor forcing and its output count

The forcing mechanism adapts [1, Lemmas 8.19–8.20, pp. 80–86]. The permissible-matching construction below makes its geometry valid for every $d \geq 2$.

This section verifies the deterministic part of the proposed local forcing argument in every $d \geq 2$. Its probabilistic inputs are kept separate: the low passage record must come from the unselected exact-shift kernel, and completion must preserve the stated interior and exterior records. No conditional success estimate or kernel-completion identity is inferred merely from the geometry below.

27.1 The corridor implication is independent of dimension

Lemma 27.1 (A strict shortcut on every new anchor geodesic)**.** *Let D be a proper geodesic metric on $\mathbb{R}^d$ inducing its Euclidean topology. Let K be a closed Euclidean ball and let $a_0, a_1 \notin K$. Write ℓ_j for the infimum of old lengths from a_j to ∂K before entering K, attained at x_j, and suppose*

$$D(a_0, a_1) \geq \ell_0 + \ell_1 + g, \qquad g > 0.$$

Let $D_f = e^{\xi f} \cdot D$, where f is continuous and supported in $\operatorname{int} K$. Fix $C > \delta > 0$, $B, e, d_0 > 0$, with

$$B + 2e < g, \qquad 4Ce < \delta d_0. \tag{27.1}$$

Assume the following concrete conditions.

1. *There is an anchor competitor of new length at most $\ell_0 + \ell_1 + B$.*
2. *Every first-to-last boundary portion of a path with new length at most B stays in an open network $O \cup T_0 \cup N \cup T_1$. Its only incidences form the cycle $O - T_0 - N - T_1 - O$. The full exterior of K and its exterior boundary layer lie in O, and $f \geq 0$ there. Every point of the outer entrance set $O \cap T_j$ has an old connection to x_j of length at most e.*
3. *There is an open $\Omega \supset N$ and points $q_j \in \Omega$ such that $d_0 = e^{\xi b} D(q_0, q_1; \Omega)$ for one real b. Every point of the terminal entrance set $N \cap T_j$ has a connection to q_j inside Ω of $e^{\xi b} D$-cost at most e. On $N \setminus (T_0 \cup T_1)$ one has $f \geq b$.*
4. *Every complete passage through T_j, in either direction and of new cost at most B, has a point z with $D_f(z, q_j; W) \leq e$, for one fixed bounded open neighborhood W.*
5. *The completed candidate passage rule has upper comparison C against D_f in W, permits concatenation, and retains a passage record from q_0 to q_1, compactly in W, of cost at most $(C - \delta) d_0$.*

Then every new geodesic between a_0 and a_1 contains ordered points $z_0, z_1 \in W$ for which

$$F_W^f(z_0, z_1) < C D_f(z_0, z_1).$$

Approximate exterior minimizers are allowed by reserving additional slack in (27.1).

Proof. The competitor is shorter than the old anchor distance, so a new geodesic must enter K. Its exterior prefix and suffix cost at least ℓ_0, ℓ_1, hence its first-to-last boundary portion costs at most B. The first corridor entered is T_0: otherwise the prefix up to the first T_1 entrance set lies in O, where old length is no larger than new length. Joining to x_1 at cost e and following the old exterior approach gives an old anchor connection of length at most $\ell_0 + \ell_1 + B + e$, contradicting the strict cost inequality. The last corridor is T_1 by the reversed argument. Also $D(x_0, x_1) \geq g$. A portion in O connecting its two entrance sets would imply $D(x_0, x_1) \leq B + 2e < g$. The network itinerary therefore contains a complete $T_0 - N - T_1$ passage. This assertion follows from the finite incidence graph, using a finite subdivision subordinate to its open cover; it allows repeated visits and exits from K.

Choose the two access points from condition 4. Between them, take the subpath after its last T_0 contact before its first subsequent T_1 contact. It stays in N, avoids the corridors in its interior, and has new length at least $d_0 - 2e$ by condition 3. As a subpath of a geodesic this gives $D_f(z_0, z_1) \geq d_0 - 2e$. The two access joins and the retained passage record have candidate cost at most

$$2Ce + (C - \delta) d_0 < C(d_0 - 2e),$$

which proves the strict inequality. Strict margins absorb arbitrarily small infimum and boundary-approximation errors. □

27.2 Two corridors in every dimension

Lemma 27.2 (Two disjoint corridors with a permissible matching)**.** *Let $d \geq 2$. Fix a closed round ball K, a closed round hole H and a closed rectangular preservation region V, with disjoint closures strictly inside K. Choose two disjoint outer caps and two disjoint terminal caps at opposite faces of V, all with fixed positive clearances. Stops lie just inside the outer caps and just outside the terminal faces. There are two disjoint simple polygonal arcs in* $\operatorname{int} K \setminus (H \cup V)$ *joining the outer stops to the terminal stops in one of the two matchings. Their short end portions inside the prescribed caps are their only permitted cap contacts.*

The arcs admit nested regular tubes and neighborhoods with exactly the incidences $O - T_0 - N - T_1 - O$, where N is a small neighborhood of V and O contains the exterior of K and a full boundary layer. There are transverse marker slabs in the tubes and a compact wall band separating a smaller network neighborhood from its thickening. Both orientations of every low passage record are included in the finite recipe catalogue.

Proof. For $d \geq 3$, the shell with the interior rectangle removed is connected and open. Begin with polygonal routes between the prescribed stops. Perturb their finitely many interior vertices inside small balls with closures in this open region. Generic choices make the arcs simple and disjoint, since two nonincident one-dimensional segments generically do not intersect in these dimensions. This gives either prescribed matching, while keeping the ends inside their caps.

For $d = 2$, choose disjoint thin routing boundary layers just inside ∂K and just outside ∂V, avoiding H. Use a parallel circle in the first layer and an offset rounded rectangle in the second. Short transverse segments join the two prescribed stops to two marked points on each of these simple closed curves. On each curve choose a base point in one of the two open arcs between its marked points. The region between the layers, with H removed, is connected. Join the base points by a simple polygonal connector whose interior avoids the layers, except for short transverse ends. Such a connector follows from polygonal connectivity and loop erasure, first between the facing layer boundaries and then across the layers at the chosen base points.

The two long sides of a sufficiently thin regular strip about the connector are disjoint. At each end, continue them along the two opposite directions of the corresponding parallel boundary curve until they reach its two marked points, then append the transverse cap segments. Because the base point lies between the marked points, these continuations are disjoint. The two strip sides determine one matching; neither matching is prescribed in advance. Small polygonal approximations inside the open clearance neighborhoods preserve disjointness and all permitted end contacts.

In both cases, the compact arcs have positive clearance from each other, the forbidden regions and the opposite caps outside their allowed end portions. Choose O and N with boundary layers narrower than the routing layers, then successively thinner tubes. This leaves exactly the stated incidences. A marker slab is a positive-width subcylinder away from a tube's ends. A wall band lies between two nested network neighborhoods inside K and away from H; it includes every interior complementary boundary. Leaving the larger neighborhood must cross that band through a positive Euclidean displacement.

A low passage record may be reversed, together with all its ordered subranges, because the graph walks and their limits are symmetric. Reversal preserves its cost, reference denominator and strict saving. It therefore supplies the permissible matching in the planar construction. The catalogue is finite, so the minimum of its positive clearances is positive. All subsequent tube and

wall choices can use this common clearance. □

No property of a Euclidean polygonal arc's D-length is used here. Later travel occurs through a connected open neighborhood of the arc, whose internal D-diameter on the compact core is finite. This distinction avoids presuming that deterministic Euclidean curves are reference-rectifiable.

Proposition 27.3 (Finite plateau recipes from a low passage record)**.** *Fix $C > \delta > 0$ and finitely many low-passage record bins with preservation boxes $V \Subset \Omega \Subset \Omega_0$ and terminal caps. Suppose an exact constant-shift sample, at one value b in a finite deterministic grid, supplies endpoints $q_j \in V$ and a passage record internal to V such that*

$$\operatorname{cost}(q_0, q_1) \le (C - \delta) d_0, \qquad d_0 = e^{\xi b} D(q_0, q_1; \Omega) \in [a'u, A'u].$$

Suppose its constant-shifted reference connection inside V costs at most $A_1 u$, and terminal cap connections cost at most $e = \epsilon u$, with the safe stop-to-endpoint connections costing at most $e/3$. Assume the unmodified reference shell gives exterior crossing lower bound at least $a > 0$, and choose $u > 0, \epsilon > 0$ so that

$$(A_1 + 4\epsilon)u < a, \qquad 4C\epsilon < \delta a'.$$

After deterministic reference-modulus truncations, choose a finite outer-cap grid having old access costs at most e, and safe attachment-to-stop costs at most $e/3$.

There is a finite family of deterministic smooth shifts, each supported in the shell, zero near its central hole and outer boundary, and exactly equal to its assigned b on the complete preservation patch V, with the following property. On an event consisting only of finitely many reference crossing, modulus, and internal-diameter bounds, every admissible exterior attachment pair and every successful low bin has one recipe satisfying all geometric conditions of Lemma 27.1, with $B = (A_1 + 2\epsilon)u$. The reference truncation event can have probability arbitrarily close to one under the fixed reference law. Recipe amplitudes and critical Cameron–Martin energies have one common finite bound.

Proof. For each of the finitely many cap/bin/grid configurations use Lemma 27.2, reversing the low passage record when needed to match its terminal caps. Both orientations are included before the finite recipe list is fixed. The outer attachments are Euclidean separated by a deterministic amount on the reference modulus event: otherwise their old distance would be less than the lower bound a. Thus disjoint outer caps can be used. Choose compact polygonal core arcs Γ_j through the two stops of each tube, away from $O \cup N$.

For all wall and marker bands, truncate the positive reference crossing distances below by one $m_* > 0$. Truncate a common local reference access modulus by a deterministic function $\omega(t) \downarrow 0$, so points within Euclidean distance t of a core have reference access to it, in a fixed buffered neighborhood, at cost at most $\omega(t)$. These are finitely many fixed local metric tests. Positivity and continuity permit the truncations at an arbitrarily small probability loss.

Choose a smooth baseline equal to zero on safe outer caps and near the outer boundary, equal to b on V and the safe terminal caps, and equal to a height H on neighborhoods of the wall and marker bands. It can be chosen nonnegative on O, at least b on N, and at most H everywhere. Take $H \ge \max(0, \max b)$ with $e^{\xi H} m_* > B$. Choose $\rho > 0$ small enough that the core ρ-neighborhoods avoid O, N, and the wall band, and $e^{\xi H}\omega(\rho) < e/3$.

Inside the core $\rho/2$-neighborhood choose a connected open domain C_j containing Γ_j compactly. Its internal reference diameter on Γ_j is finite almost surely: internal distances in an open connected subset of a length space with Euclidean topology are finite and continuous, so compactness applies.

Truncate these finitely many diameters above by one $K_* < \infty$. Only now choose J_0 so large that $e^{-\xi J_0} K_* < e/3$ and $-J_0 \le \min(0, \min b)$. Lower the baseline to $-J_0$ on neighborhoods of the closures of C_j, interpolating inside the core ρ-neighborhoods. This leaves $f = b$ on V, $f \ge b$ on N, and $f \ge 0$ on O. Safe cap costs can only decrease.

A path of cost at most B cannot cross a wall band, so it remains in the thickened network; partial wall excursions are included in that thickening. A complete tube passage avoiding the core ρ-neighborhood would cross its marker slab entirely at weight $e^{\xi H}$, costing more than B. At a point where it enters the core neighborhood, join to Γ_j at cost $e/3$, travel inside C_j to the terminal stop at cost $e/3$, and use the safe terminal cap at cost $e/3$. This gives the access condition in Lemma 27.1.

The competitor follows the two unchanged exterior approaches, the safe outer caps, the depressed cores, the safe terminal caps, and the retained reference connection in V. Its new length is at most $\ell_0 + \ell_1 + 2e + A_1 u$. The middle portion outside the corridors has $f \ge b$ and lies inside Ω, which gives the old constant-shifted denominator lower bound required by that lemma. The stipulated inequalities are exactly its two strict margins.

All choices followed a deterministic order: finite geometry; crossing and modulus cutoffs; H; ρ and C_j; diameter cutoff K_*; then J_0. Thus the recipes are fixed before any old geodesic or successful sample selects one. Finiteness gives common amplitude and $\dot{H}^{d/2}$ energy bounds. The latter is invariant under Euclidean translation and dilation. □

The denominator is deliberately not retained as a new metric readout. Only the passage record restriction internal to V is retained from the exact-shift sample. The number $e^{\xi b} D(q_0, q_1; \Omega)$ is source information used in its selection and estimates. Since the recipe need not equal b on Ω, it is not identified with $D_f(q_0, q_1; \Omega)$. The lower bound on the new central passage follows instead from $f \ge b$ there. This resolves the larger-boundary layer issue without any candidate Weyl rule.

27.3 Plateau recipes with local white-noise scores

For odd d, a general critical Sobolev shift has a nonlocal Gaussian score. The plateau construction can be made inside a smaller class whose scores are local in the white noise. Suppose $\kappa \ge 0$, $\int \kappa > 0$, and write $K_t(x) = t^{-d} \kappa(x/t) / \int \kappa$. Let $g \in C_c^\infty$ have all desired plateau values on slightly larger neighborhoods than those used above. Choose a smooth probability density η supported in $[a, 2a]$, with a smaller than every deterministic clearance, and set

$$f(x) = \int_a^{2a} \eta(t) (K_t * g)(x)\, dt, \qquad v(y,t) = \frac{\eta(t)}{\int \kappa} t^{-(d-1)/2} g(y).$$

Shifting the space–scale white noise by v adds precisely f to the logarithmic field. Positivity and unit integral of K_t preserve all plateau identities and required one-sided inequalities; support stays inside the shell. The score is $\langle W, v \rangle$, and

$$\|v\|_2^2 = \frac{\|g\|_2^2}{(\int \kappa)^2} \int_a^{2a} \eta(t)^2 t^{-(d-1)}\, dt < \infty.$$

It uses only a buffered local white-noise region. Under a physical translation/dilation, the shift $v_{z,r}(y,t) = r^{-(d+1)/2} v((y-z)/r, t/r)$ has the same norm. The density calculation for completed passage arrays still requires the stated conditional-kernel interface; this construction supplies the local Gaussian score and its uniform energy bound.

27.4 Shrinking-support stability

The stability mechanism is adapted from [1, Lemma 8.22, pp. 86–88].

Lemma 27.4 (A saturated fixed-anchor ratio is stable)**.** *Suppose D is proper and has a unique geodesic P between fixed anchors, parametrized by fractions of its length. Suppose $P([t_0,t_1]) \Subset U$ and*

$$F_U(P(t_0),P(t_1)) = C(t_1-t_0)D(a_0,a_1),$$

with old upper comparison C. Let f_n range over shifts with one bounded amplitude, supported in balls of radii tending uniformly to zero and centered near a fixed compact middle subarc of $P((t_0,t_1))$. Assume every new completion has upper comparison C against D_{f_n} and preserves old full passage record restrictions outside a fixed-factor enlargement of the support. Assume ordered subpath capture and local endpoint joining for those records. Then, uniformly over all such shifts and completions,

$$D_{f_n}(a_0,a_1) \longrightarrow D(a_0,a_1), \qquad \sup_t D(P^{f_n}(t),P(t)) \longrightarrow 0,$$

and

$$\frac{F_U^{f_n}(P^{f_n}(t_0),P^{f_n}(t_1))}{(t_1-t_0)D_{f_n}(a_0,a_1)} \geq C - o(1).$$

If corridor forcing supplies a strict shortcut in the shrinking middle neighborhood, the displayed ratio is also strictly below C for all sufficiently small supports.

Proof. There are fixed $0 < m \leq 1 \leq M$ such that $mD \leq D_{f_n} \leq MD$. Let ω_n be the maximum old diameter of the enlarged support balls. Uniform continuity on the compact middle neighborhood gives $\omega_n \to 0$. Cut an old or new minimizing path at its first and last support visits. An old connection between these cut points costs at most ω_n, and its new cost at most $M\omega_n$. This proves that old and new anchor distances differ by $O(\omega_n)$. The new minimizing segment spanning all visits has old length at most $m^{-1}M\omega_n$; thus every old and new prefix length differs by $O(\omega_n)$. Properness confines the minimizing paths to one compact set, and their fractional parametrizations are uniformly D-Lipschitz. Every subsequential limit is the old minimizing path with the old fractional parametrization. Uniqueness proves uniform convergence, also uniformly over support locations, by a contradiction-subsequence argument.

Take a new passage record between the new fractional endpoints. On actual realizations cut at first and last visits to a slightly larger ball lying outside the removed exterior patch. Ordered capture retains old exterior prefix and suffix records with no increase of their combined cost. Join the two cut points using the old upper comparison at cost $C\omega_n + o(1)$ and correct the endpoints at cost $2C\eta_n + o(1)$, where $\eta_n = \sup_t D(P^{f_n}(t),P(t))$. These joins stay in U because the changed middle subarc has positive clearance and reference short connections stay nearby. If the passage record avoids the changed neighborhood, only endpoint corrections are needed. Old saturation therefore gives

$$C(t_1-t_0)D(a_0,a_1) \leq F_U^{f_n}(P^{f_n}(t_0),P^{f_n}(t_1)) + C(\omega_n + 2\eta_n).$$

The positive denominator converges, proving the lower ratio estimate. The changed neighborhood is separated from the old outer geodesic portions; uniform path convergence puts the new shortcut strictly between the same fractional times. Following the remaining new geodesic at upper factor C gives the strict upper ratio bound. □

27.5 The counting argument has no planar input

The weighted measure argument is the one in [1, Lemma 8.17 and Remark 8.18, p. 80].

Suppose a finite candidate family has at most cN Euclidean radii. At each radius its centers form a fixed mesh with bounded overlap, and at each center/radius at most J labels and recipes can satisfy the output alignment test. Assign a deterministic weight w_β. Let an output event B_β require that its own fixed-anchor reference geodesic traverses a designated unchanged shell and that the output's internal shell crossing is at least κw_β. Also include an output length cutoff L_*. Charging each weight to arclength inside its enlarged ball yields pointwise

$$\sum_\beta w_\beta \mathbf{1}_{B_\beta} \le \kappa^{-1} J C_d(cN) L_*. \tag{27.2}$$

The shell lower bound must be part of B_β itself. Traversal or label alignment alone is insufficient for this charge.

For comparison, if every point of an old middle geodesic arc lies in at least $\beta_0 N$ good inner holes, and a hole's first-to-last outer boundary visits account for at most $A w_\beta$ old geodesic length, integration along the arc gives

$$\sum_{\text{hit good candidates}} w_\beta \ge A^{-1}\beta_0 N \operatorname{len}_D(P_{\text{mid}}).$$

This is a source-side statement. It uses minimizing reference paths and the surrounding connection bound, not candidate noncollapse.

Finally, let R be the same measurable output ratio for every candidate and every N, and put $H_N = \{C - \epsilon_N \le R < C\}$, $\epsilon_N \downarrow 0$. If fixed-recipe shift and completion provide output submeasures with

$$\sum_\beta w_\beta \nu_\beta(1) \ge aN, \quad \nu_\beta \le \Lambda\mu, \quad \nu_\beta(H_N \cap B_\beta) = \nu_\beta(1),$$

then (27.2) implies $aN \le \Lambda K N \mu(H_N)$. But $\mathbf{1}_{H_N} \to 0$ pointwise, including on $\{R = C\}$, so $\mu(H_N) \to 0$, a contradiction. The hypotheses about source selection, likelihood clipping, exact kernel completion, and all-point candidate coverage must be verified before applying this count. Geometry alone does not provide them.

28 Restricted passage-law rigidity under explicit reference inputs

This section identifies the graph passage law using the established strong reference and local cores of Section 2. The probability-transfer, spatial-amplification and corridor arguments concern the graph records and were proved above.

28.1 The precise interface

Hypothesis 28.1 (Established strong reference and local cores)**.** Fix $d \ge 2$, $0 < \gamma < \sqrt{2d}$, $Q = d/\gamma + \gamma/2$, and $\xi > 0$. Put $\alpha = \xi/\gamma$. There is one Borel assignment $g \mapsto D_g$ for the logarithmic field and its continuous perturbations with the following properties. It is a continuous length metric inducing the Euclidean topology. For the actual pinned logarithmic field used in the proof it is proper and geodesic; the same then holds for every bounded continuous shift used below, by its uniform metric comparison. No properness assertion is made for arbitrary unbounded

continuous perturbations. Its internal metrics are given by consistent local Borel rules. It satisfies the simultaneous continuous Weyl rule and, for fixed $z \in \mathbb{R}^d$, $r > 0$, the exact coordinate rule

$$D_g(z + rx, z + ry) = D_{g(z+r\,\cdot)+Q\log r}(x, y). \tag{R1}$$

The reference has a positive local Euclidean upper Hölder exponent. It satisfies the countable deterministic-anchor property of Lemma 22.3: every nonconstant ambient geodesic segment compactly contained in an open W contains a nonconstant subsegment of a unique geodesic between two fixed-catalogue anchors, and that entire anchor geodesic is compactly contained in W.

These properties are supplied by the previously established results in Section 2 and the dimension-specific source checks in Lemma 22.3. They are stated here to make the passage-law theorem's interface explicit. They are not new unproved assumptions about the reference. If the controlled arrays were first constructed for the finite-range field, Proposition 21.2 is applied before this theorem: it constructs the actual canonical pinned-field arrays in the same deterministic units and preserves the local kernel and comparison interfaces. The global deterministic-anchor argument below is run under that canonical pinned law, where properness is available. Its intrinsic conclusion can afterwards be transferred back locally to the finite-range field.

Lemma 28.2 (Established uniqueness for fixed deterministic endpoints)**.** *For the reference metric, almost surely each fixed deterministic pair of distinct points has a unique geodesic. The assertion holds simultaneously for a countable collection of pairs and after a countable collection of fixed Cameron–Martin shifts.*

The first assertion is [2, Theorem 1.2] for $d \geq 3$ and [5, Theorem 1.2] for $d = 2$. Countable intersections and equivalence under fixed Gaussian shifts give the stated simultaneity; a constant change of pin rescales the metric and leaves its minimizing paths unchanged.

Hypothesis 28.3 (Full passage records)**.** A passage law, coupled with the actual field and actual reference internal metrics, has the following properties. All assertions are made on a countable library of buffered charts, open ports, rational allowances, and fixed affine coordinate maps; the library is enlarged before extraction whenever another fixed test is needed.

1. A finite-total-cost passage records endpoints, a compact range, a finite nonnegative cost bound, and finite ordered subranges with additive subcost bounds. Infinite-cost compactification points are auxiliary records and are not called passage records below. Every passage record compactly inside an open chart has actual discrete realizations with arbitrarily small strict spatial and cost losses. Conversely every sequence of bounded-cost actual paths in a fixed compact set has a captured subsequential passage record. Ordered subdivision at prescribed cost increments, extra marked points, and cutting at buffered testing faces are retained. Spatial mesh and normalized edge overshoot tend to zero. All statements hold jointly for finite lists and then by a countable diagonal. Open-margin restriction compatibility is imposed; no continuity of restriction to a closed chart boundary is assumed.

2. Before any comparison with D_g, a uniform local endpoint joining modulus is available in the controlled families used below. Passage records meeting at a common endpoint can be concatenated in any prescribed open neighborhood with arbitrarily small total cost loss. The infimal compactly confined cost $F_U(x, y)$ is symmetric, has zero diagonal, is domain monotone, and obeys the triangle inequality. Finite local chains give finite costs in connected open domains.

3. On separated buffered patches the restrictions have product conditional kernels depending only on the actual local field. They are compatible with a full Borel kernel. The same kernels apply after every fixed locally equivalent Gaussian field change. The entire fixed white-noise background and finite local recipe scores may be retained, as in Proposition 23.5; passage kernels still depend on the field, not on extra noise conditional on that field.

4. With $v_i = b^{-i}$ and $c_i = \gamma^{-1} \log v_i$, the indexed kernel is exactly

$$K_i(g) = \left(S_{v_i^{-\alpha}}\right)_* K(g + c_i), \tag{R2}$$

where S_a multiplies every cost by a. A physical restriction on $z + rU$, divided by v_i^α, has kernel K_i at $g(z + r\,\cdot) + Q \log r - c_i$. These are transported-experiment identities, not candidate Weyl rules. The full countable collection of fixed offsets is retained together.

5. The controlled indexed laws used in the proof have jointly extractable full closed record collections and actual continuous reference internal-metric readouts on compact exhaustions. A further extracted law has a diagonal realization by the actual discrete arrays, with all fixed offsets, vanishing meshes, and the endpoint joining bounds retained. Witnesses and event indicators can be adjoined. At such extractions numerical upper and lower comparisons survive. The latter follows from the first two items and the retained internal readouts as explained immediately below.

For the discrete arrays constructed in this note, the last item is provided, after the initial finite upper comparison, by Lemma 20.6. That lemma first fixes each selected index, then chooses a sufficiently late actual resolution matching both its marked laws and its calibration moments. It uses the upper comparison to obtain common reference moduli and finite-port moment bounds, and does not assume a positive lower comparison or regular variation.

An upper comparison C means that every compactly confined D_g-rectifiable path can be followed, in every prescribed open neighborhood of its range and between prescribed endpoint ports, at candidate cost at most C times its reference length plus any positive allowance. A lower comparison c means that every passage record of cost a and endpoints x, y satisfies $a \geq cD_g(x, y)$. These are simultaneous almost-sure statements.

For clarity about the last interface item, lower comparisons pass through attainability and continuous endpoint readouts. For an upper comparison, subdivide a limiting reference path into finitely many pieces lying in smaller buffered open charts. Approximate their endpoints by the countable port library. The retained reference internal metrics bound the preceding laws' connection costs in those charts, and their upper comparisons supply passage records with the corresponding bounds. Capture and the common endpoint modulus pass these finite follows to the limit. Refining the subdivision and allowances gives the claimed upper comparison. Thus this survival does not presuppose candidate noncollapse or a candidate Weyl rule.

Theorem 28.4 (Restricted passage-law rigidity)**.** *Assume Hypotheses 28.1 and 28.3. Suppose the indexed candidate internal diameters of a surrounding sphere in a fixed buffered annulus have uniformly tight upper tails. Suppose also that the finite-anchor guarded calibration of Lemma 28.6 is retained and is nonzero with positive probability. Then there is a deterministic $C \in (0, \infty)$ such that, almost surely,*

$$F_U(x, y) = CD_g(x, y; U) \tag{R3}$$

on every open chart, simultaneously under compact exhaustion. The finite-total-cost full records compactly confined in U are precisely the ordered records realizable by reference-rectifiable curves

in U whose subcurve lengths, multiplied by C, are bounded by the recorded subcosts. In particular these intrinsic records are determined by the field. Raw auxiliary coordinates attached to a closed chart's boundary are not part of this assertion.

28.2 Path calculus and annular comparison

Lemma 28.5 (Intrinsic path calculus)**.** *An upper comparison $C < \infty$ gives $F_U \le C D_g(\cdot,\cdot;U)$ and local endpoint continuity of F_U. If $c > 0$ is a lower comparison, every finite-total-cost passage record compactly confined in U has a reference-rectifiable realization in U with the same full range and traversal order, and with each subcurve length at most its subcost divided by c. Consequently its cost is at least $cD_g(x,y;U)$.*

Proof. Follow a near-minimizing reference path, shrink its endpoint ports, and use capture and endpoint joining. The triangle inequality gives

$$|F_U(x,y) - F_U(x',y')| \le C\{D_g(x,x';U) + D_g(y,y';U)\}.$$

Short reference connections on a compact subset remain in U, so the right side tends to zero with the endpoints.

For the lower assertion, parametrize actual realizing walks by their normalized edge cost. Mark the first vertex after every rational parameter time, retaining the vanishing edge overshoot. A diagonal capture gives marks x_t satisfying $D_g(x_s,x_t) \le |t-s|/c$. They extend to a reference-Lipschitz curve of length at most the recorded cost divided by c. To check that its image is the entire retained range, take any range point y, choose realizing vertices tending to y, and extract their cost times $t_n \to t$. Retaining this extra mark gives $D_g(y,x_s) \le |t-s|/c$ for every rational s, hence $y = x_t$. This argument may be made separately for each y; it does not require uncountably many simultaneous selections. Zero-cost pieces have zero reference diameter and are singletons. Apply the same construction to every ordered subpiece and concatenate. □

Lemma 28.6 (The finite anchor calibration excludes zero upper constant)**.** *Retain the finite guarded anchor observable F_{I_n}/b_n jointly with the local passage records, where $b_n \to \infty$. Suppose its limit Z is nonzero with positive probability, compactly confined passage records are attainable by actual graph walks, and local endpoint joining is available. If the passage law has a finite deterministic upper comparison with D, its optimal upper constant is strictly positive.*

Proof. If the optimal constant were zero, intersect the probability-one comparisons at constants $1/k$. For each of the finitely many anchor pairs choose a finite-length reference connection compactly inside its prescribed open trace domain. The zero upper comparison supplies attainable passage records of arbitrarily small cost in that same domain. Correct their endpoint errors by local joining. Actual graph walks give continuous cell-contact traces, and vanishing Euclidean cell mesh preserves strict confinement. The short portions from the deterministic anchor points to their containing original cells add only a bounded number of cells and stay inside the same open domain for all sufficiently large n. Inserting the finite guard set adds at most its cardinality to the distinct-cell cover of each trace. Thus every summand of F_{I_n}/b_n is bounded above in the joint limit by an arbitrarily small passage record cost. There are only finitely many summands, and all finite insertion and endpoint errors vanish since $b_n \to \infty$. This forces $Z = 0$, a contradiction. □

Lemma 28.7 (Finite ball chaining)**.** *Consider finitely many closed Euclidean balls of diameter at most 2ε, whose designated interior holes cover a compact path range. Ball B has weight $w_B > 0$. Suppose a completed hole-to-exterior traversal pays a charge at least aw_B, while any two points*

on ∂B can be joined in its prescribed buffered annular neighborhood at cost at most Aw_B. One can then connect points within 4ε of the original endpoints at total joining cost at most A/a times the sum of the traversal charges. This statement holds for $d \geq 2$ and for captured passage records whenever the positive charges are available on their actual marked realizations.

Proof. For a continuous path this is Lemma 25.1, with the same charges and endpoint error. For a retained compact range, fix a finite cover by open holes. It covers all sufficiently late actual realizing paths. Lemma 25.2 gives a uniform charge $a'w_B$ for every actual traversal, for each $0 < a' < a$, before the number of cuts is counted. Its positive minimum bounds that number. Apply the charged-ball lemma to the actual paths, capture the resulting bounded list of marks, and let $a' \uparrow a$. The joins remain in the prescribed open neighborhood for fine covers. □

Beginning of Theorem 28.4. Uniform indexed candidate annular joining tightness and positivity of the fixed reference annular crossing give constants $M < \infty$ and $m > 0$ such that the joint event

$$\sup_{u,v \in \partial B_2} F_{A_{1,3}}(u, v) \leq M, \qquad D_g(\partial B_1, \partial B_2; A_{1/2,3}) \geq m$$

has probability above the threshold of Proposition 24.4 at every index. Both tests read only the observed annulus; neither reads its central hole. This distinction is essential for separated-annulus Gaussian comparison. The joining tightness is an instance of compact internal-diameter tightness in a connected open annulus. Apply that proposition on the admissible small-chart atlas and a summable sequence of label windows, using the fixed-density local transfer in Remark 24.6. Each late ancestry and all its radius-8 history spheres then fit in one admissible observation window; the global pin is not silently moved through an increasing exhaustion. The exact reference rule (R1) and candidate identity (R2) give the same physical quantum weight to the upper cost bound and lower charge. The resulting arbitrarily fine covers use balls of radius $2r$ and central holes of radius $r/128$. Every completed hole-to-exterior traversal pays the reference charge, and boundary contacts can be joined through the tested candidate annulus. Lemma 28.7, followed by the preexisting endpoint joining modulus to remove its vanishing endpoint errors, yields a finite deterministic upper comparison. Whole-ball diameter events would not justify this application, since their nested observation regions overlap. The nonzero calibration and Lemma 28.6 exclude optimal upper constant zero. Write $C \in (0, \infty)$ for the optimal upper constant and $c \in [0, C]$ for the optimal lower constant. The optimal constants themselves are valid: intersect the countable comparisons approaching their infimum or supremum and then send allowances to zero.

These numerical constants are unchanged by a fixed index shift. Indeed, after expressing a positive lower comparison intrinsically by Lemma 28.5, both comparison properties are local tests. Fixed local Gaussian equivalence transfers them in both directions under $g \mapsto g + c_i$. Reference Weyl scaling gives $D_{g+c_i} = v_i^{\alpha} D_g$, which cancels the candidate cost factor in (R2). For $c = 0$ the lower comparison is simply nonnegativity. This argument concerns local field restrictions; it does not claim that adding a constant preserves the law of a globally pinned field. The exact coordinate rule transports the same numerical comparisons through the actual affine atlas. □

28.3 High and low witnesses at most quantum labels

This amplification and the following equality-tangent argument adapt [1, Propositions 8.14–8.15, pp. 76–79].

Lemma 28.8 (Amplification of comparison witnesses)**.** *Suppose $0 \leq c < C < \infty$. There are a fixed closed annular band $K \Subset W$, with W an open annular boundary layer, and $p_H > 0$, such that*

for every $\delta \in (0, C)$, more than three quarters of the labels in every sufficiently late window $[N, 2N]$ have probability at least p_H of a nonconstant reference geodesic P, minimizing in W, whose entire range lies in K and whose endpoints lie on its two different boundary components, with

$$F_W(P(0), P(1)) \geq (C - \delta) \operatorname{len}_{D_g}(P). \tag{R4}$$

The geometry and p_H do not depend on δ. For every $c_1 \in (c, C)$ there are $p_L > 0$ and finitely many fixed closed rectangular boxes R with buffered boundary layers R^+ such that more than three quarters of the labels in each sufficiently late window have probability at least p_L of a passage record confined in one such R, with endpoints on its designated opposite faces, and

$$a < c_1 D_g(x, y; R^+). \tag{R5}$$

Fixed positive lower and finite upper denominator cutoffs and compact reference moduli can be added at a fixed loss of probability.

Proof. Choose a reference regularity event of probability arbitrarily close to one on which the separated band crossing costs at least $a_0 > 0$ and all surrounding-boundary connections cost at most $A_0 < \infty$ in their prescribed bands. These choices use only the fixed reference metric, precede δ, and apply to every label.

If (R4) had probability below a sufficiently small fixed p_H at at least one quarter of arbitrarily late labels, its complementary event, intersected with reference regularity, would pass the all-point annulus threshold. Follow a compact reference geodesic and stop upon successive exits from the good covering balls. In each completed traversal retain the subarc from the last inner-boundary visit before the first outer-boundary visit. Its whole range lies in the corresponding closed band K, and it minimizes in the buffered boundary layer W. These disjoint geodesic subarcs have length at least $a_0 w$ in the physical quantum units and, on the complementary event, admit candidate follows saving at least $\delta a_0 w$ relative to factor C. Follow the remaining pieces at factor C. Lemma 28.7, now using reference surrounding joins, gives

$$\operatorname{len}_{D_g}(P) \leq A_0 \sum w + o(1).$$

Thus the new candidate cost bound is at most

$$C \operatorname{len}_{D_g}(P) - \delta a_0 \sum w + o(1).$$

After finite realizations and endpoint corrections, and then refinement of the cover, this improves the upper factor to any number greater than $C - \delta a_0 / A_0$. Short geodesic subdivision and the buffered atlas extend the improvement to every compactly confined reference path, contradicting optimality of C.

Here the compact-band restriction is imposed in the definition of (R4) before estimating its probability. No assertion is made that cutting an arbitrary high-ratio witness preserves its high ratio: the absence of the restricted event itself supplies the saving for each traversal subarc above. Also every witness to this restricted event has length at most $\operatorname{diam}_{D_g(\cdot,\cdot;W)} K$. This finite random variable has the same law at every label. Together with the positive band-crossing cutoff, it gives fixed length and compact-modulus cutoffs uniformly as $\delta \downarrow 0$. These are the cutoffs used later in Lemma 28.9; that later extraction lemma is not an input to the present amplification proof.

For the low witness, cover a middle annular region by central quarters of finitely many small axis-parallel cubes $q + [-a, a]^d$. Every hole-to-exterior path visits such a central quarter and subsequently exits the cube while still in the testing region. According to its first exit direction,

retain the piece from its last visit to $x_j = q_j \pm a/2$ before its first visit to $x_j = q_j \pm a$. It lies in a rectangular slab with width $a/2$ and other side lengths $2a$. The $2d$ directions for each of the finitely many cubes give the required finite box list. Their enlarged boundary layers fit strictly in the observation annulus. On a fixed high-probability reference event all the opposite-face denominators are at least a common $a_0 > 0$.

Suppose (R5) had probability below a sufficiently small p_L on at least one quarter of arbitrarily late labels. The transferred good annuli then give lower cost $c_1 D_g(x, y; R^+)$ to every such rectangular passage record. In each completed traversal of an actual realizing path mark one slab piece. The remaining pieces retain the old lower comparison. Relative to that old bound each marked piece pays an additional $(c_1 - c)a_0 w$. For $c > 0$ use the intrinsic form from Lemma 28.5; for $c = 0$ use nonnegativity. At a fixed finite cover the positive marked charge bounds the number of traversals even when $c = 0$, so finite ordered capture is valid. Reference triangle inequalities and the finite ball chain give, for a passage record from x to y,

$$a \geq c D_g(x, y) + (c_1 - c)a_0 \sum w - o(1), \qquad D_g(x, y) \leq A_0 \sum w + o(1).$$

This improves the lower comparison to every constant smaller than $c + (c_1 - c)a_0/A_0$, a contradiction. Work first with a threshold strictly between c and c_1 to retain strict slack after the stated cutoffs. Since the reference field law is fixed across indices, its denominator and modulus truncations have a common arbitrarily small failure probability. This proves all assertions. □

28.4 A saturated fixed-anchor segment in a further law

Lemma 28.9 (Equality tangent and fixed-anchor saturation)**.** *Under the preceding hypotheses with $c < C$, a further controlled indexed extraction has upper constant exactly C, has a confined low passage record with a deterministic strict margin below C with positive probability, and, separately with positive probability, has deterministic anchors $a_0 \neq a_1$, rational $0 < t_0 < t_1 < 1$, and a bounded deterministic open U such that its unique reference geodesic P satisfies*

$$P([t_0, t_1]) \Subset U, \qquad F_U(P(t_0), P(t_1)) = C(t_1 - t_0) D_g(a_0, a_1). \tag{R6}$$

Proof. Choose $c < c_1 < c_2 < C$ and $\delta_k \downarrow 0$. The high and low favourable label sets in Lemma 28.8 intersect in more than half the labels of every sufficiently late window. Choose one deterministic label tending to infinity from each such intersection. Adjoin measurable high and low witnesses, with dummy values off their respective events, and retain their event indicators. Closed record spaces and compact curve spaces provide measurable selections after fixed cutoffs; completed universally measurable selections can equivalently be replaced by Borel versions under the fixed law.

High witnesses have range in one fixed compact annular band inside W. Truncate their reference lengths to $[a, A] \subset (0, \infty)$ at a fixed loss of positive probability. Their fractional parametrizations have a common tight Euclidean modulus, since they are A-Lipschitz in the reference metric and the latter has the Euclidean topology. Retain the actual reference internal metrics on a compact exhaustion. A joint controlled extraction, realized by Lemma 20.6, gives a positive-length limiting path P_*, and the identities

$$D_g(P_*(s), P_*(t); W) = |s - t| L$$

follow by passage through those readouts. Thus it is an internal reference geodesic.

If its endpoints had a passage record in W costing strictly less than CL, that passage record's range would have positive clearance in one compact subchart. Full-record attainability in the

preceding laws and the common endpoint joining bound would give passage records violating their high inequalities (R4) for large k. The limiting geodesic is therefore saturated. Every interior subsegment is saturated too: a cheaper passage record for it, joined to upper-comparison follows of the complementary pieces, would contradict saturation of the whole path.

For the low witness pass to one fixed box in the finite list. Retain its actual denominator and fixed lower and upper cutoffs. Closed range confinement, opposite-face endpoints, and continuous denominator readouts pass to a passage record with ratio at most $c_2 < C$. The retained indicators ensure that the high and low events each retain positive probability, without asserting their simultaneous occurrence. The inherited upper comparison is at most C, while saturation excludes a smaller one. A lower comparison exceeding c_2 would contradict the low witness through Lemma 28.5.

Keep the actual whole-space reference law pinned outside these testing patches. Choose a short interior saturated segment whose length is smaller than its ambient distance to W^c; a shorter ambient competitor could not leave W, so it is an ambient geodesic. Lemma 22.3 applies to this actual fixed reference law in every dimension $d \geq 2$. Proposition 22.4, including its countable choice of anchors, rational fractions and rational-ball domain, now gives (R6). The shared core is selected from the current reference geodesic. No local seed-code measurability or uniqueness for arbitrary interior endpoint pairs is needed in this deduction. □

28.5 A high-probability local forcing catalogue

Lemma 28.10 (The complete catalogue used in rigidity)**.** *Suppose favourable finer indices have a confined rectangular passage record of cost at most $(C-\delta)D_g(q_0,q_1;\Omega_0)$ with probability at least $\beta > 0$, for fixed $\delta > 0$ and finitely many geometric types. For every $\theta > 0$, there exist a fixed base power $M = b^m > 1$ and a finite deterministic list of localized white-noise recipes such that the outer indexed experiment at volume Mv_i has a local field-and-auxiliary-record event of probability at least $1-\theta$ with the following property. Every reference geodesic making a deep hit of the central hole admits one recipe and one previously sampled auxiliary passage record for which Lemma 27.1 applies. The list has uniformly bounded amplitude and noise energy, preserves the full real-field passage record patch exactly, and permits subsequent likelihood clipping by Lemma 24.7 for Proposition 23.8. The good event and auxiliary sampling are specified before any old geodesic or old candidate record is selected.*

Proof. Reduce the strict margin slightly. Truncate the source denominator below by $a' > 0$, choose $\epsilon > 0$ with $4C\epsilon < \delta a'$, and use reference moduli to bin the two endpoints into finitely many small face ports. Enlarge the closed box to an open preservation patch $V \Subset \Omega \Subset \Omega_0$. Each terminal cap lies inside Ω, contains its port and a deterministic stop just outside V, and has constant-shifted reference access cost at most ϵ; the safe stop-to-endpoint part costs at most $\epsilon/3$. These are finite geometric bins. Shrinking the denominator domain from Ω_0 to Ω only increases its internal distance, so preserves the low-ratio inequality.

Truncate the denominator above and the internal reference connection in V below a finite cost bound A_1, retaining a fixed positive success probability for the union of bins. All these choices precede M. Apply Proposition 26.1 to place finitely many separated trial windows in the shell and amplify this source event. That proposition supplies independent fine-field opportunities, a deterministic finite exact-shift grid, and the full auxiliary arrays. Its conditional coarse profiles are used only for the probability estimate; the resulting success event is a local field-and-array event.

Truncate the outer reference crossing below by $a > 0$ and its surrounding connection and

modulus bounds above. Choose $M = b^m$ so large that

$$(A_1 + 4\epsilon)M^{-\alpha} < a.$$

In outer units put $u = M^{-\alpha}$, $e = \epsilon u$, and $B = (A_1 + 2\epsilon)u$. The exact transported input on a trial chart ψ_j is

$$g \circ \psi_j + Q \log \rho_j + \gamma^{-1} \log M + s.$$

The finite compensation grid in the spatial amplification is chosen only after M and the trial charts are fixed. Rounding changes the sampled input field law by a bounded constant; it never assumes persistence of a particular passage record under a field change. Retain a successful passage record at the exact grid value where it was sampled. Its source denominator is in $[a'u, A'u]$, and its reference connection through V costs at most $A_1 u$. The actual outer recipe has plateau s on V, not $\gamma^{-1} \log M + s$. Indeed the latter constant belongs to the transported source experiment; multiplying its reference metric and all its candidate costs by $M^{-\alpha}$ cancels exactly the factor $e^{\xi\gamma^{-1}\log M} = M^{\alpha}$.

For every terminal bin, outer attachment-cap pair, trial, and exact shift, Proposition 27.3 constructs the disjoint tubes, walls, marker bands, safe caps, and depressed cores. The two inequalities required there are precisely the displayed choice of M and $4C\epsilon < \delta a'$. The outer caps are finite by the truncated reference modulus, so the recipe list is finite and its guarantee holds for every admissible attachment pair. The auxiliary denominator on Ω is used only as source information; only the passage record internal to V is retained on completion. The pointwise lower bound on the recipe outside the tubes supplies the central denominator lower bound in the new metric.

Choose every recipe from the buffered plateau convolutions of Lemma 22.5. They preserve the exact values on V, walls, cores, and unchanged boundary layers, as well as all required one-sided bounds. Their finite common noise energy is invariant under physical translation and dilation. Add the finitely many reference wall crossings, core internal-diameter bounds, and access moduli used by the tube lemma, allocating each a sufficiently small failure-probability bound. The total loss is below θ. These tests belong to the unaugmented local field-and-auxiliary record. After all-point transfer is applied, use Lemma 24.7 to add the finite likelihood clips without changing the previously chosen failure threshold or geometric scale sequence. □

28.6 The fixed-law output contradiction

The weighted contradiction follows [1, Lemma 8.17, Remark 8.18 and Proposition 8.28]; the preceding lemmas verify its present model-specific inputs.

Completion of Theorem 28.4. Suppose $c < C$. Pass to the further law of Lemma 28.9. It has the same optimal upper constant C, a strict low witness, and a positive-probability event (R6) for one fixed-anchor geodesic. Apply Lemma 28.8 to this law to obtain favourable finer labels of positive density in every sufficiently late window. The preceding catalogue and Proposition 24.4, with exact-shift auxiliary kernels but without score augmentation, then provide arbitrarily fine good annuli at these labels. Choose its failure threshold before constructing the catalogue. After the finite recipes are fixed, Lemma 24.7 clips all their likelihood scores while deleting at most half the successes per ancestry. Absorb this factor into the success density.

Restrict the source saturation event to a positive-probability event S_* on which the saturated interval and a compact middle subinterval have reference lengths bounded below, the entire anchor length is at most $L_* < \infty$, and all needed compact containment and clearance bounds hold. Fix a

deterministic compact set containing that middle subarc. Keep this same whole-space Gaussian law and its exterior pin throughout; no change of reference law is made after choosing the global fixed-anchor event.

In the catalogue good event also require that every two points on the surrounding sphere have a reference connection, inside a fixed buffered compact shell, of cost at most Aw. This is an internal diameter bound, not a minimum crossing bound. Also require a lower deep-hit crossing aw and a lower crossing κw in a further unchanged shell. Here

$$w = (Mv_i)^\alpha$$

is the deterministic weight of the candidate label. All crossing readouts are internal distances between fixed testing faces in buffered bands. Thus they are local and charge ambient traversals by last-inner/first-outer clipping. Their high-probability cutoffs can be chosen before applying all-point transfer.

With probability tending to one, every point of the old middle arc lies in at least $\beta_0 N$ good central holes. A ball hit by that arc accounts for at most Aw old geodesic length: the first-to-last surrounding-boundary subpath minimizes, and its boundary endpoints have the stipulated reference join. Integrating along the arc gives the pointwise source bound

$$\sum_{\text{hit good balls}} w \geq A^{-1}\beta_0 N \operatorname{len}_{D_g}(P_{\text{mid}}). \tag{R7}$$

At each hit use one fixed deterministic ordering to select the first successful trial, bin, attachment caps, and recipe. This partitions the source events without changing their total weight.

We now impose the source-side stability cutoff before defining output submeasures. For each N the possible physical centres, radii, labels, and recipes form a finite deterministic family of shifts f_β and noise directions u_β. Over those whose balls deeply hit the compact middle arc, let ω_N be the largest old reference diameter of an enlarged support ball, and let η_N be the largest value of

$$\sup_{t\in[0,1]} D_g(P^{f_\beta}(t), P(t)) + |D_{g+f_\beta}(a_0, a_1) - D_g(a_0, a_1)|.$$

An empty maximum is zero. The shifts across all N form a countable deterministic list, vanish near the fixed pin, and have bounded amplitude. Gaussian equivalence gives unique anchor geodesics for all of them simultaneously; their measurable fixed-anchor readouts make these maxima measurable. Lemma 27.4 gives $\omega_N + \eta_N \to 0$ on the source cutoffs, uniformly over all permissible passage completions. Add to the error the indicator of failure of the required new compact confinement or fractional interval placement; that indicator is eventually zero by the same lemma. We obtain an error $E_N \to 0$ almost surely, depending on the old reference field but not on the completion. Choose a deterministic $\varepsilon_N \downarrow 0$ with $\mathbb{P}[E_N > \varepsilon_N] \to 0$.

On the intersection of S_*, the covering event, and $\{E_N \leq \varepsilon_N\}$, (R7) is bounded below by a fixed positive multiple of N. This event has probability $\mathbb{P}(S_*) - o(1) > 0$. Hence the expected retained source weight is at least $a_* N$ for one $a_* > 0$. This uses a pointwise bound on the retained event; no uniform-integrability assertion is needed for weights discarded by the stability cutoff.

For each deterministic candidate recipe, first perform its fixed noise shift and complete the full passage object according to Proposition 23.8; only afterwards restrict to its selected source event. The likelihood clipping gives output submeasures

$$\nu_\beta \leq \Lambda\mu$$

with one finite Λ and the same full field-and-passage law μ for every N, β. The corridor implication supplies a strict shortcut on every new anchor geodesic. Shrinking-support stability puts it in the same fixed fractional interval and gives

$$C - \varepsilon_N \le R := \frac{F_U(P(t_0), P(t_1))}{(t_1 - t_0) D_g(a_0, a_1)} < C. \tag{R8}$$

Here R is one and the same measurable function of the output field and output passage object. Define it as $+\infty$ off the required endpoint domain. Local endpoint continuity and a countable compact exhaustion give a Borel version of F_U, and unique fixed-anchor geodesics have measurable versions. Thus no candidate-dependent ratio is being substituted into (R8).

Define B_β from output data: its own fixed-anchor geodesic traverses the further unchanged shell; the output internal shell crossing is at least κw_β; its anchor length obeys the common output cutoff; and its unchanged outer-average alignment holds. The old anchors lie outside all sufficiently small candidate buffers, and corridor forcing makes the new geodesic enter the modified compartment, so selected outputs belong to B_β. The normalization average and the further shell are unchanged by the recipe. These are therefore valid output tests, including the shell lower charge itself.

At each Euclidean radius the centre mesh has bounded overlap. For a fixed centre and radius, bounded alignment allows only a bounded number of labels; the recipe list is finite. There are $O(N)$ radii. Charging each $w_\beta \mathbf{1}_{B_\beta}$ to the output geodesic's arclength inside its enlarged ball yields exactly the pointwise multiplicity estimate (27.2):

$$\sum_\beta w_\beta \mathbf{1}_{B_\beta} \le K_* N.$$

The output length cutoff follows from the source cutoff and bounded recipe amplitude; including it in every B_β makes the bound pointwise for all outputs.

Set $H_N = \{C - \varepsilon_N \le R < C\}$. The source weight, domination, and support properties now give

$$a_* N \le \sum_\beta w_\beta \nu_\beta(1) = \sum_\beta w_\beta \nu_\beta(H_N \cap B_\beta) \le \Lambda K_* N \mu(H_N).$$

But $\mathbf{1}_{H_N} \to 0$ pointwise, including when $R = C$, so $\mu(H_N) \to 0$ by bounded convergence. This is a contradiction. Thus $c = C > 0$ for the original law.

Lemma 28.5 upgrades the lower comparison to internal distances; the upper comparison gives the reverse inequality, proving (R3). It also realizes every retained finite-total-cost ordered record as a reference-rectifiable curve with the asserted subcost bounds. Conversely, given such a curve compactly in U, subdivide its prescribed subarcs finely and follow them at factor C in shrinking open neighborhoods. Correct consecutive endpoint errors by the local joining modulus. The actual subranges stay close to the desired ranges and meet increasingly fine sets of points on them, so their ordered ranges converge in Hausdorff distance. Capture and closedness give the desired passage record. Countable exhaustion gives the simultaneous intrinsic statement. □

Remark 28.11 (The reference input and the graph conclusion). The proof assumes no candidate Weyl rule, candidate field measurability or candidate positivity. Those properties follow after the passage law is identified. The strong reference and its common deterministic-anchor property are supplied by the cited results and Lemma 22.3; Proposition 22.4 is their deterministic application to a saturated graph limit. The extraction keeps this actual fixed reference law, so no uniform seed catalogue for varying weak assignments is used.

29 Propagation from selected records to every intensity

The relative-unit and first-failure strategy is adapted from [1, Proposition 8.31, Lemma 8.32, Proposition 17.2 and Theorem 17.3]. The proof below supplies the Poisson-intensity implementation.

This section separates the normalization argument from its remaining passage-rigidity input. The result below is an implication with explicit inputs. In particular, the scalar exponent, sampling concentration, and Gaussian kernel lemmas alone do not supply its restricted-rigidity or physical-label-transfer hypotheses.

Let $F_I \geq 1$ be one plus the sum of the finitely many guarded distinct-cell connection costs used in the anchor construction. Write their deterministic endpoint pairs and open trace cubes as $(a_\ell, b_\ell; U_\ell)$, $1 \leq \ell \leq m$, with $a_\ell \neq b_\ell$ in U_ℓ. The reference field in this normalization experiment is the stationary finite-range field X. Fix

$$\alpha = \xi/\gamma > 0, \qquad b(I) = (\mathbb{E}F_I^p)^{1/p}, \qquad 8p < \min\{q_0,\, 2d/(\xi\gamma)\},$$

decreasing p further if a local Gaussian comparison uses Hölder. Here q_0 is the exponent in the anchor-tail estimate. The quantile exponent and Lemma 17.4 give

$$b(I) = I^{\alpha + o(1)}. \tag{29.1}$$

For $0 < I \leq 1$, a fixed trace and the local Poisson count bound $b(I)$ above; always $b(I) \geq 1$.

The following terminology specifies the inputs. A *full marked law* retains the background field, its actual reference metric and internal readouts, the normalized F_I, and the closed path/passage records in a countable buffered atlas. It retains all fixed intensity offsets and the required Gaussian recipe scores. Upper passage records have actual discrete realizations. Bounded-cost traces admit the marked restrictions and captures used for lower bounds. Only costs and paths in strict interior buffers are used to identify a metric; closed-boundary coordinates may remain auxiliary marks. All these laws and finite experiments are placed in one separable marked-law space. We write S_c for multiplication of all passage costs by $c > 0$.

The marked compactness and transport input. For each fixed $s > 0$, $K < \infty$, the past-envelope condition

$$b(J) \leq K b(I)(J/I)^s, \qquad 1 \leq J \leq I, \tag{29.2}$$

supplies joint asymptotic compactness of the full marked laws in units $b(I)$, a common local endpoint-joining modulus, tight compact internal diameters, and uniformly integrable calibration powers up to some order strictly above $4p$. These properties extend to every fixed compact window of intensity multipliers. Closed lower witnesses and strict open upper witnesses are stable in this record topology. For the present discrete model these assertions follow in order from Theorem 9.2, Theorem 20.3, and Proposition 23.5. Indeed, integration of the anchor tail gives every calibration moment strictly below q_0; the same anchor construction supplies simultaneous adaptive endpoint joins and compact internal diameters. Its final guard removal gives these estimates for the original graph. The full closed-record construction then supplies all ordered upper-allowance records, attainability and capture, and the fixed-density closure supplies the local conditional kernels. No lower comparison or rigidity is used in any of these three steps.

For a fixed $t = e^{\gamma c}$ a smooth cutoff equal to one on the whole observation buffer gives the exact local identity

$$\mathbb{E}[(F_{tI}/b(I))^q] = \mathbb{E}[R_c(F_I/b(I))^q]. \tag{29.3}$$

The same identity holds for the entire local record experiment with the field coordinate translated back. The actual reference readouts are translated using their established smooth Weyl rule. Fixed-density kernel closure preserves these identities jointly with all retained background coordinates. They imply continuity for moving $t_n \to t > 0$ of the *whole* transformed experiment. No continuity of one cloud coordinate in a prescribed nested Poisson coupling is assumed.

Lemma 29.1 (Moving compact intensity windows)**.** *Suppose the base full marked laws in deterministic units B_n have a controlled joint limit and the stated calibration moment bounds. For deterministic $t_n \to t \in (0,\infty)$, their local experiments at intensity $t_n I_n$, still in units B_n, converge to the fixed t transported experiment. If the own-unit ratios have positive limits, the assertion also holds in their own units.*

Proof. Fix any finite projection of the atlas and choose one smooth Cameron–Martin cutoff ψ equal to one on all its observation buffers and reference internal-metric domains. Put $c_n = \gamma^{-1}\log t_n$ and $c = \gamma^{-1}\log t$. The exact finite Poisson identity expresses the transformed law as the base joint law weighted by R_{c_n}, followed by the map $X \mapsto X - c_n\psi$. Its passage records are unchanged; its retained internal reference metrics whose domains lie in the plateau are multiplied by $e^{-\xi c_n}$. The retained Gaussian background coordinates undergo the corresponding fixed Cameron–Martin translation. This is an identity of experiments, not a comparison of two clouds at nearby intensities.

On a compact interval of c the Gaussian likelihoods satisfy $\|R_{c_n} - R_c\|_{L^1} \to 0$ and have uniformly bounded moments of every fixed order. Approximate R_c in L^1 by bounded continuous cylinder functions; the approximation error is uniform in n because the background marginal is fixed. Those functions pass through joint weak convergence. The displayed field and noise translations and scalar multiplications of genuine internal reference readouts are continuous in their retained coordinates and in c_n. If the projection also retains the global ambient reference metric, its output is instead the actual Weyl metric $e^{-\xi c_n\psi}\cdot D_X$; it is not multiplied by the plateau scalar. For fixed c, this is a measurable readout of the fixed background, so Lusin approximation under that fixed background marginal passes it through the joint limit. For moving c_n, the Weyl order bound

$$e^{-\xi|c_n-c|\|\psi\|_\infty}(e^{-\xi c\psi}\cdot D_X) \le e^{-\xi c_n\psi}\cdot D_X \le e^{\xi|c_n-c|\|\psi\|_\infty}(e^{-\xi c\psi}\cdot D_X)$$

gives local uniform convergence for each fixed background. A likelihood truncation makes the corresponding approximation uniform in the weighted experiments. This uses no properness of D_X and no continuity of internalization. Thus the fixed finite projection converges to its t transport. Every finite atlas projection is treated this way; their identities are consistent because they came from the same finite Poisson experiments. Tightness and a countable determining family give the full marked assertion. Hölder preserves a slightly smaller moment order than the available one, uniformly on a compact t interval. Finally, dividing all costs by a convergent positive own-unit ratio is continuous. No plateau-scalar identity is asserted for a global ambient distance. □

The two remaining geometric inputs. We use the following precise forms, so that the logical role of an annular-transfer or rigidity proof can be checked. All transferred tests are confined to one annular observation window, for example $A_{1/16,4} = B_4 \setminus \overline{B}_{1/16}$, with their path ranges in the smaller buffer $A_{1/2,7/2}$. An upper surrounding-join test bounds the cost of joining every pair of points of ∂B_2 inside $A_{1,3}$. A lower crossing test charges every ∂B_1–∂B_2 piece in its buffered annulus. Fixed inner and outer ports for an upper crossing lie in the same observation window, strictly on opposite sides of these testing spheres. Last-inner/first-outer cutting charges any ambient traversal. No whole-ball diameter or record inside the central hole is included in a separated-annulus probability experiment.

(L) *Physical quantum-label transfer.* At each fixed base experiment, let a deterministic set of sufficiently small quantum labels have positive upper density in arbitrarily late logarithmic windows. Use its actual joint local shape, coarse-history, and recipe kernels on the corresponding physical annuli. If at those labels the passage surrounding-annulus joining tests have cost at most M, and the reference separated crossing costs at least m, with the fixed sufficiently high probability required by the annular argument, then the base passage law has upper comparison at most $C_g M/m$ with the reference metric. Conversely, passage crossing charges at least m and reference joining costs at most M give lower comparison at least $m/(C_g M)$. The constant C_g depends only on the chosen buffered geometry. Comparisons concern all confined passage records and their intrinsic restrictions. This statement requires the physical indexed kernels; matching one-dimensional crossing marginals does not suffice.

(R) *Restricted rigidity.* An admissible joint indexed passage law with a finite deterministic upper comparison to the reference metric, a nonzero finite calibration coordinate, and the controlled full-record extractions used by the rigidity argument is a deterministic multiple CD_X of the reference local passage structure. Its comparisons and fixed-index identities are preserved under those controlled extractions. This input asserts no compactness for an arbitrary preassigned moving sequence of discrete intensities and offsets.

In (L), upper comparison means the availability of realizable confined upper passage records along reference paths; lower comparison bounds the cost of every confined passage record. Fixed Gaussian changes transfer each probability-one comparison back to the appropriate pinned-field atlas. In quantum units, the reference constant Weyl rule and the exact indexed kernel identity transfer the same comparison constant to every fixed offset.

Lemma 29.2 (Two-sided comparison from annular covers)**.** *The full marked interface, the physical all-point label-transfer proposition, and the exact reference coordinate rule imply (L). Thus (L) has no further planar path-intersection input.*

Proof. The physical all-point label-transfer proposition gives arbitrarily fine finite central-hole covers outside a summable exceptional sequence. The reference coordinate rule and the actual indexed identity assign the same positive physical weight to both readouts on each ball. Apply Theorem 25.3. Its upper part gives realizable follows with factor M/m; its lower part bounds every confined record with factor m/M, including when the old lower constant is zero. The uniform prelimit shell charge is established there before the discrete cuts are counted. Fixed changes in the testing geometry give C_g. Countable compact and strict-buffer exhaustions make the comparisons simultaneous on each open domain. □

Theorem 29.3 (All-intensity propagation, conditional on (L) and (R))**.** *Assume (29.1), the marked compactness and transport input above, the physical-label transfer (L), and restricted rigidity (R). Assume the Euclidean cell mesh and the finite-guard removal errors vanish in units $b(I)$. Then every sequence $I \to \infty$ has a subsequence whose full local passage laws, divided by $b(I)$, identify with CD_X, where C is deterministic and belongs to one fixed compact subinterval of $(0,\infty)$. After transfer to the canonical pinned whole-space field h, the unrestricted hop metrics converge jointly with h and its actual reference metric, locally uniformly, to CD_h along that subsequence.*

A fixed annular median then gives deterministic normalizers $a(I)$ for which the hop metrics for the canonical pinned field converge locally uniformly in probability to the chosen reference normalization, at every intensity $I \to \infty$. These normalizers satisfy

$$a(I) = I^{-\alpha+o(1)}, \qquad \frac{a(tI)}{a(I)} \longrightarrow t^{-\alpha} \qquad (t>0).$$

Properness is used only for this canonical reference law. No properness claim for an arbitrary continuous perturbation or for the stationary finite-range auxiliary field is needed.

29.1 Calibration for the finite port observable

Lemma 29.4 (Positive comparison and fixed scalar bounds)**.** *Suppose a full marked limit in units $B_n \to \infty$ retains $F_{I_n}/B_n \Rightarrow Z$ with $\mathbb{E}Z^p = 1$. If it has a finite deterministic upper comparison to the reference metric, its optimal upper comparison constant is positive. For a rigid law CD_X there are fixed reference variables $0 < W_- \le W_+ < \infty$, with $\mathbb{E}W_+^p < \infty$, such that*

$$CW_- \le Z \le CW_+. \tag{29.4}$$

Consequently

$$C_- := \|W_+\|_p^{-1} \le C \le C_+ := \|W_-\|_p^{-1}, \qquad R_0 := C_+/C_- < \infty. \tag{29.5}$$

Proof. If the optimal upper comparison were zero, intersect its probability-one assertions with constants $1/k$, $k \in \mathbb{N}$. For every one of the finitely many port pairs choose a finite-length reference path compactly contained in its prescribed cube. Upper comparison gives confined passage records of arbitrarily small total cost. Realize them in the actual unguarded clouds and attach the endpoints using the common joining modulus. A contact trace of a graph walk is covered by its own cells, up to a fixed endpoint correction. Inserting the finite guard set increases a distinct-cell cover by at most the number of inserted sites. Thus the normalized sum of the guarded port costs is arbitrarily small; the additive one in F_I also vanishes. This forces $Z = 0$, a contradiction.

For each trace cube choose a fixed larger connected open buffer U_ℓ^+ containing its closure and all vanishing mesh corrections. For the lower bracket take $W_- = \sum_{\ell=1}^m D_X(a_\ell, b_\ell; U_\ell^+)$. Delete the finite permanent guards. Rootless screening gives a tight total deletion overhead, hence a vanishing error in the chosen units. Each bounded-cost port trace has a captured local passage record. Its intrinsic lower comparison in U_ℓ^+ gives the corresponding summand of CW_-. For the upper bracket take $W_+ = \sum_{\ell=1}^m D_X(a_\ell, b_\ell; U_\ell)$, or a larger sum in fixed slightly smaller connected trace domains containing the ports. Near-minimizing reference paths are compactly confined in those open domains. Realizable upper passage records followed by guard insertion prove the upper inequality. Both brackets are therefore read from fixed internal reference metrics; no ambient-distance functional for the auxiliary field X needs to be transported under a compact-cutoff change of law.

The endpoints are distinct, so $W_- > 0$ almost surely. Connectedness and the length topology give $W_+ < \infty$ almost surely. The primary internal-diameter estimate has moments of every order below $2d/(\xi\gamma)$ on fixed geometry. A local smooth Gaussian correction and reference Weyl scaling transfer a sufficiently small positive moment to the field X. Our choice of p leaves room for the required Hölder inequalities; no all-order moment at a fixed endpoint is used. Taking pth moments in (29.4) proves (29.5). □

By (29.3), the record moment bound of order $4p$, Hölder, and the inverse density comparison, for every compact $T \Subset (0, \infty)$ one has at past-envelope bases

$$0 < c_T \le \frac{b(tI)}{b(I)} \le C_T < \infty \quad (t \in T). \tag{29.6}$$

For clarity, with $U = F_I/b(I)$ we have $\mathbb{E}U^p = 1$ and $\mathbb{E}U^{2p} \le C$. Choose $\delta > 0$ so that

$\mathbb{P}(R_c < \delta) < 1/(4C)$ uniformly on the compact c interval. Then $\mathbb{E}[R_c U^p] \geq \delta/2$ by Cauchy–Schwarz. The upper bound is immediate from Hölder. Also

$$|\mathbb{E}[(R_c - R_{c'})U^p]| \leq \|R_c - R_{c'}\|_2 \|U^p\|_2 \leq C_T|c - c'|.$$

Thus the deterministic profiles in (29.6) are equicontinuous in $\log t$. They have positive continuous subsequential limits, without yet being powers.

If a base and its fixed offsets have been identified as rigid laws, the fixed-index identity transports the same scalar in units $t^\alpha b(I_n)$. Applying the port bracket and its moment calibration at both indices gives

$$R_0^{-1} t^\alpha \leq \lim_{n\to\infty} \frac{b(tI_n)}{b(I_n)} \leq R_0 t^\alpha. \tag{29.7}$$

This bound is sufficient for propagation; equality is not assumed. More precisely, if $r(t) = \lim b(tI_n)/b(I_n)$ and C_t is the own-unit scalar at the fixed multiplier t, the exact indexed identity gives $C_t r(t) = C_1 t^\alpha$. Thus for any fixed $u, v > 0$,

$$R_0^{-1}(u/v)^\alpha \leq r(u)/r(v) \leq R_0(u/v)^\alpha. \tag{29.8}$$

The constant is R_0, not a product of two independent profile bounds: the common base scalar cancels. Uniform integrability of the retained calibration at both offsets is what places both C_u and C_v in the same interval.

29.2 Relative units in a backward array

The following consequence of (L) is useful because it needs no compactness uniform in quantum units at its outset.

Lemma 29.5 (Relative-unit comparison)**.** *Let G_i, $i \geq 0$, be actual own-unit passage laws, and let $\beta_i > 0$. Suppose that for every fixed base i its relative offset law at quantum volume $2^{-\ell}$, in that base's quantum units, is*

$$(S_{\beta_{i+\ell}/\beta_i})_\# G_{i+\ell}. \tag{29.9}$$

These are the actual local kernels on physical annuli. Suppose each G_i has the same sufficiently high-probability tests on one fixed buffered shell: separated passage and reference crossings cost at least $m > 0$; passage and reference surrounding connections cost at most $M < \infty$; and each has a crossing of cost at most M between fixed ports whose closures force the separated crossings. Every upper join and upper port crossing is confined to the annular window just specified, and every lower test is a closed witness in its buffered annular subregion. Suppose the fixed-base realization and capture operations and (L) hold. Then some deterministic $T < \infty$ satisfies

$$T^{-2} \leq \beta_i/\beta_j \leq T^2 \quad (i, j \geq 0).$$

Proof. Take T strictly larger than $C_g(M/m)^2$, with additional fixed slack for the port buffers. Fix a base i. If, in arbitrarily late windows $\ell \in [N, 2N]$, at least one quarter of the labels satisfy $\beta_{i+\ell} > T\beta_i$, then (29.9) makes the passage crossing charge there at least mT, while reference joining remains bounded by M. Apply the lower half of (L) to this deterministic label set inside the fixed base experiment. It gives the intrinsic lower comparison $mT/(C_g M)$ with the reference metric. The base shell tests have positive probability of a passage crossing of cost at most M and reference port separation at least m. This contradicts $m^2 T/(C_g M) > M$.

If at least one quarter of arbitrarily late labels instead satisfy $\beta_{i+\ell} < \beta_i/T$, passage joining costs at most M/T while the reference crossing charge remains at least m. The upper half of (L) gives comparison at most $C_g M/(mT)$. Apply it to the base reference crossing of length at most M; strictly nested ports and actual realizability preserve a separated passage crossing. Its cost is below m, contradicting the base passage lower test. Increase T once to absorb all fixed margins.

Thus for each fixed i and all sufficiently large N, fewer than one quarter of the indices in $[i+N, i+2N]$ have $\beta_k > T\beta_i$, and fewer than one quarter have $\beta_k < \beta_i/T$. Fix i, j and take N large for both. Their windows overlap in $N - O(|i-j|)$ indices. Removing the first upper-exception set and the second lower-exception set leaves a common k, so $\beta_j/T \le \beta_k \le T\beta_i$. Interchange i, j for the reverse inequality. No starting value of N uniform in i was required. □

29.3 Strict success tests near the rigid laws

Let $\mathcal{R}$ be the class of admissible rigid normalized marked laws with scalar in $[C_-, C_+]$. Their actual reference readouts are retained. The marked closure input makes $\mathcal{R}$ compact: the scalar has a convergent subsequence, strict interior upper passage records pass by approximation, and compact closed lower witnesses preserve the matching lower bound. Boundary auxiliary coordinates are included in their compact extraction rather than identified with an unproved boundary-distance functional.

Fix $s = \alpha/2$. Choose $R > 1$ so that the bounds in (29.7) for $t \in [1/2, 2]$ lie strictly inside (R^{-1}, R). Choose an integer m_0 with

$$R_0 2^{-m_0(\alpha - s)} < 1/4. \tag{29.10}$$

Cone records moved to their nearest lower dyadic level have the fixed past-envelope constant $K_{\mathrm{start}} = 4R_0$; this is verified below. Put

$$C_{\mathrm{run}} = (R2^s)^{m_0 - 1}, \qquad K_* \ge C_{\mathrm{run}} \max\{R, K_{\mathrm{start}}\}.$$

These constants are fixed before any first-failure sequence is chosen. Let $\mathcal{P}$ be the compact class of limiting full marked laws of bases satisfying (29.2) with K_*, enlarged by their factor-two windows and by $\mathcal{R}$. The compactness input gives this compact class and its common upper joining modulus; it provides no lower control on all of $\mathcal{P}$.

Lemma 29.6 (A strict shell-test neighborhood in the full record topology)**.** *For every prescribed sufficiently small failure-probability bound $\varepsilon > 0$ there are $m > 0$, $M < \infty$, and an ambient open neighborhood $\mathcal{O} \supset \mathcal{R}$ such that every law in $\overline{\mathcal{O}} \cap \mathcal{P}$ satisfies all shell tests of Lemma 29.5 with probability at least $1 - \varepsilon$. This statement permits zero-cost connected ranges in members of $\mathcal{P}$ and makes no identification of raw records on the boundary of a chart.*

Proof. Choose one compact annular chart K_{obs} whose interior contains $\overline{A}_{1/2,7/2}$ and which lies in the annular observation window. All tests below use its record coordinate. In particular, the closed crossing range $K_0 = \{1 \le |x| \le 2\}$ lies strictly inside this chart. Fix closed inner and outer ports strictly on opposite sides of the two testing spheres and contained in $A_{1/2,7/2}$.

Let $\mathcal{B}_m$ be the event that there exists a one-piece record whose endpoints lie on $\partial B_1, \partial B_2$, whose range is contained in K_0, and whose total cost is at most m. The set of possible witnesses is compact: endpoints and ranges are compact, and their costs belong to $[0, m]$. Each constraint is closed. Consequently $\mathcal{B}_m$ is closed in the full state space, and $\nu \mapsto \nu(\mathcal{B}_m)$ is upper semicontinuous. Every rigid law has positive factor at least C_-; its intrinsic lower comparison bounds this event

by the lower tail of one positive reference annular crossing. Thus m can be chosen so that its probability is uniformly much smaller than ε on $\mathcal{R}$.

For upper tests take a finite $\delta/4$-net $z_1, \dots, z_k$ of ∂B_2. For each pair i, j ask for a record with endpoints in the open $\delta/4$-ports about z_i, z_j, range compactly in $A_{1,3}$, and cost strictly below $M/2$. Ask for one more such open-port connection across the two fixed testing ports, confined in $A_{1/2,7/2}$. Each is an open hit condition on a closed record collection: endpoint conditions are open, $R \Subset A_{1,3}$ means the compact range belongs to the open subset $\{R \in \mathcal{H}(K_{\mathrm{obs}}) : R \subset A_{1,3}\}$, and the cost bound is strict. Their finite intersection $\mathcal{E}$ is therefore open. The reference upper surrounding-join and port costs are finite almost surely. Since $C \le C_+$, choose M so large that the rigid laws satisfy these tests with substantially stronger cost and probability margins, uniformly in the eventual finite net. One reference all-pairs surrounding-join bound supplies all the net tests at once; no union bound depending on the net size is needed.

Now choose δ from the common local joining modulus on $\mathcal{P}$, so that outside a probability bound smaller than $\varepsilon/5$, every two endpoints in the relevant compact annular neighborhood at Euclidean distance at most δ can be joined inside $A_{1,3}$ at cost at most $M/8$. On that event, $\mathcal{E}$ supplies all surrounding-boundary joins at cost below $M/2 + 2(M/8) < M$. All reference tests can be expressed by strict inequalities on finitely many retained continuous internal-metric readouts (including their compact suprema or infima); their joint event $\mathcal{E}_D$ is open and has uniformly high probability. Decrease m and increase M if necessary before choosing the net.

With small fixed numerical error allocations, let $\mathcal{G}$ be the set of laws for which $\nu(\mathcal{B}_m)$ is below its allocated probability bound and both $\nu(\mathcal{E})$ and $\nu(\mathcal{E}_D)$ exceed their allocated success probabilities. Portmanteau makes $\mathcal{G}$ open, and the strict margins give $\mathcal{R} \subset \mathcal{G}$. Every law in $\mathcal{G} \cap \mathcal{P}$ has all the desired shell tests, by the preceding modulus bound and a finite union bound.

It is not enough simply to take $\mathcal{O} = \mathcal{G}$: a closed bad-witness probability can jump upward at a limit. Instead fix a metric on the marked-law space and use compactness. The disjoint compact sets $\mathcal{R}$ and $\mathcal{P} \setminus \mathcal{G}$ have positive distance η (if the latter is empty, take any $\eta > 0$). Set $\mathcal{O} = \{\nu : \mathrm{dist}(\nu, \mathcal{R}) < \eta/3\}$. Then $\overline{\mathcal{O}} \cap \mathcal{P} \subset \mathcal{G}$, as required.

The strict interior location of K_0 prevents unidentified chart boundary coordinates from entering a lower witness. A zero-cost range which traverses the shell is also detected: perform its first and last cuts on actual realizing walks, retain adjacent vertices, and capture the ordered zero-cost subrecord. Thus no path parametrization or positive lower metric bound is hidden in the closed-test construction. Finally, $\mathbb{E}Z^p = 1$ for the calibration coordinate is not treated as an open weak-topology condition. It passes in the controlled class by its uniform integrability; the finite-port bracket supplies the same property on the rigid class. If a fixed local Gaussian density is needed, choose the initial error allocation through its fixed modulus of absolute continuity before this construction. □

Write $A_j = b(2^j)$ and let ν_j be its full own-unit law. For $j \ge m_0$ call j *successful* when

$$\nu_j \in \mathcal{O}, \qquad R^{-1} < \frac{b(t2^j)}{A_j} < R \quad (1/2 \le t \le 2), \qquad \frac{A_{j-m_0}}{A_j} < 2^{-m_0 s}. \tag{29.11}$$

The full multiplier interval is needed for a real-intensity past envelope. Equicontinuity of the deterministic profiles makes it a legitimate uniform numerical test. Success is a property of a deterministic probability law and deterministic normalizing constants. Hence all first-failure indices below are deterministic; no scale is selected by conditioning on a particular field or cloud realization.

29.4 Initialization at cone records

Lemma 17.1, or the same lemma using (29.1), supplies $I_n \to \infty$ and $\epsilon_n \downarrow 0$ such that for every $J \geq 1$,

$$\frac{b(J)}{b(I_n)} \leq \max\{(J/I_n)^{\alpha-\epsilon_n}, (J/I_n)^{\alpha+\epsilon_n}\}. \tag{29.12}$$

Retain each fixed offset vI_n in units $v^\alpha b(I_n)$. Its past-envelope multiplier differs by at most $e^{\epsilon_n|\log v|}$, which tends to one. Taking the discrete limit first for each fixed v gives one common upper joining modulus for all the indexed laws. Finite deterministic chains in $A_{1,3}$ turn this modulus and its tight finite connection costs into uniform surrounding-boundary join tests, measurable from the annular records alone. Reference crossing charges have a positive fixed lower quantile. The upper half of (L) therefore initializes a finite deterministic upper comparison. Lemma 29.4 excludes zero, and (R) identifies the resulting law as CD_X.

Any subsequent selected offsets v_k are realized by a controlled diagonal of the actual discrete arrays: choose $n(k)$ so late that $\epsilon_{n(k)}|\log v_k| \to 0$, $|\log v_k|/\log I_{n(k)} \to 0$, $v_k I_{n(k)} \to \infty$, and $v_k^\alpha b(I_{n(k)}) \to \infty$, and approximate the required finite lists of records and moduli to a vanishing tolerance. All fixed extra offsets are added successively. More precisely, after the common finite upper comparison has been initialized, Lemma 20.6 bounds all indexed limit moduli by the actual reference internal metrics and all indexed calibrations by CW_+. Fixed-offset higher moments allow the late actual representatives to match these moment bounds as well as their record laws. This supplies the controlled extractions required in (R) before rigidity is proved; it is not an assertion about arbitrary preassigned offset/resolution sequences.

Let $r_n = \lfloor \log_2 I_n \rfloor$. Extract $2^{r_n}/I_n \to t \in [1/2, 1]$. Moving-window transport gives a rigid own-unit law at 2^{r_n} and at every fixed finite window of dyadic levels around it. The ratio at $t \in [1/2, 1]$ is bounded away from zero by (29.7). For $J \leq 2^{r_n}$, (29.12) and $\alpha - \epsilon_n \geq s$ then give

$$b(J) \leq K_{\text{start}} A_{r_n} (J/2^{r_n})^s$$

Indeed, uniform profile convergence gives $b(tI_n)/b(I_n) \geq (2R_0)^{-1} t^\alpha$ on $[1/2, 1]$ for all large n, whereas the cone coefficient at that dyadic base is $t^{\alpha-\epsilon_n}$. Their quotient is at most $2R_0 t^{-\epsilon_n} \leq 4R_0$ when $\epsilon_n \leq 1$. This proves the asserted initial envelope with a constant fixed before $\mathcal{P}$ and $\mathcal{O}$ were chosen. Openness of $\mathcal{O}$, uniform convergence of the multiplier profiles, and the strict choices of R, m_0 imply that every fixed finite forward window from these starts is eventually successful. A diagonal subsequence yields starts with $r_n, r_n + 1, \ldots, r_n + n$ all successful.

29.5 A successful run has one fixed past envelope

Suppose every level from r through j is successful, and the start r has its K_{start} past envelope. For $r \leq k \leq j$, write $j - k = am_0 + b$, $0 \leq b < m_0$. Iterating the last and middle tests of (29.11) gives

$$\frac{A_k}{A_j} \leq 2^{-am_0 s} R^b \leq C_{\text{run}} 2^{-s(j-k)}.$$

For $J < 2^r$, combine the initial real-intensity envelope with this bound on A_r/A_j. For $2^r \leq J \leq 2^j$, choose $2^k \leq J \leq 2^{k+1}$ and use $b(J) \leq RA_k$. In both cases the fixed choice of K_* yields

$$b(J) \leq K_* A_j (J/2^j)^s, \qquad 1 \leq J \leq 2^j. \tag{29.13}$$

There is no dependence on the length of the successful run.

29.6 Exclusion of a late first failure

Suppose unsuccessful levels occur arbitrarily late. From each long start r_n choose the first subsequent failure f_n, and put $j_n = f_n - 1$. Then $j_n - r_n \to \infty$. Equation (29.13) supplies compactness, the common joining modulus, and calibration uniform integrability in units A_{j_n}. The retained calibration has pth moment one.

For every fixed $i \geq 0$, the level $j_n - i$ is successful for all large n. Its ratio to A_{j_n} lies in $[R^{-i}, R^i]$. Retain positive finite limits $r_i = \lim_n A_{j_n-i}/A_{j_n}$, $r_0 = 1$, and its own-unit law G_i, jointly with every fixed offset. Each G_i belongs to $\overline{\mathcal{O}} \cap \mathcal{P}$ and so has the same two-sided shell tests. Put $\beta_i = r_i/2^{-i\alpha}$. The actual discrete array gives, at base i and offset $2^{-\ell}$, exactly the limiting quantum-unit law $S_{\beta_{i+\ell}/\beta_i} G_{i+\ell}$. Thus Lemma 29.5 gives

$$T^{-2} \leq \beta_i \leq T^2 \quad (i \geq 0). \tag{29.14}$$

Multiplying the common own-unit joining bounds by these bounded factors gives uniform upper annular tests in backward quantum units. Compact factor-two windows fill between their dyadic labels using the fixed-density transport. The upper half of (L) now supplies a finite upper comparison for the base-0 law. Its unit calibration moment excludes zero. Once this comparison is obtained, the exact fixed-offset kernel and the reference constant Weyl rule transfer it to every fixed offset, with the same constant in quantum units.

Varying-index extractions needed by (R) are now supplied by Lemma 20.6. First retain the bounded relative factors in (29.14) and the common own-unit moduli. The finite upper comparison just obtained gives the uniform indexed reference bounds in that lemma. Select each desired index before choosing its original discrete representative; choose that representative late enough to match the finite record list, vanishing errors, and calibration moments, with diverging physical intensity and cost unit. No comparison from (R) was needed to obtain this extraction. Hence (R) makes the preceding law rigid.

The fixed forward-window transport gives compactness and uniform integrability for intensities through 2^{j_n+1}, still in preceding units. The failed scale's normalizer has a positive finite limiting ratio, bounded already by the successful scale's factor-two test. For clarity, no estimate requiring its success is used: its past envelope already follows from the preceding one with the fixed constant $\max\{K_* R 2^s, R^2\, 2^s\}$, by splitting past intensities at 2^{j_n}. Alternatively its compact-window laws and moments follow directly from the preceding Gaussian density. The available moment order is strictly above $4p$; Hölder with conjugate exponent arbitrarily close to one preserves a moment order above $4p$ after this fixed-window change. Its own-unit limit is a normalized fixed-offset version of the preceding rigid law, so it belongs to $\mathcal{R}$. Its law test therefore passes eventually. Uniform convergence of its multiplier profiles and (29.7) make its entire window test strict. All its fixed offsets are rigid by transport from the preceding base, with their own calibration moments retained. Applying the pairwise form (29.8) to $u = 2^{1-m_0}$ and $v = 2$ gives, without two independent R_0 losses,

$$\limsup_n \frac{A_{f_n - m_0}}{A_{f_n}} \leq R_0 2^{-m_0\alpha} < 2^{-m_0 s}$$

by (29.10). Its last test passes too. This contradicts first failure. Every sufficiently large dyadic level is successful.

29.7 All intensities, global distances, and the median

Choose an original cone start beyond the final failed level. All subsequent levels form one successful run, so the original constant K_* in (29.13) applies at every later dyadic level. Restarting from an

arbitrary successful level could enlarge that constant and is not needed. For any dyadic sequence, repeat the backward-array argument inside the same compact class $\mathcal{P}$: its fixed backward laws have the shell tests, relative units are bounded, and (L) and (R) identify the limit. For arbitrary $I_n \to \infty$, write $I_n = 2^{j_n} t_n$ with $t_n \in [1, 2)$ and extract $t_n \to t$. Moving-window transport gives the same deterministic-multiple identification. Its scalar belongs to (29.5).

First transfer the intrinsic local identification, using the countable buffered atlas and fixed Gaussian densities, to the canonical pinned field h. This uses local restrictions only. Upper passage records along a compact reference path can be glued through finitely many atlas charts by the common endpoint-joining modulus. For the lower bound, use finite triple buffers $K_a \Subset V_a \Subset W_a$ inside the prescribed trace domain, with the inner interiors covering its compact trace region. Charge first exits from V_a after a start in K_a using the comparison and reference distance in W_a. The captured range lies in $\overline{V}_a \Subset W_a$, so no comparison on a closed chart boundary is used. The minimum of these finitely many reference exit distances is positive. A bounded-cost trace therefore has only finitely many completed exits; capture them with ordered marks, retain the final incomplete piece, and sum the intrinsic local lower bounds. Thus the transferred identification holds on every compact trace region, with the same deterministic scalar C.

The transferred intrinsic comparisons, realizable upper records, vanishing cell mesh and common local joining modulus meet Lemmas 31.1 and 31.2. They confine all bounded-cost first-exit prefixes and identify unrestricted distances locally uniformly along the extracted subsequence. The probability localization in Section 31 applies with the fixed positive lower scalar C_-.

Properness of the canonical reference law is the established input recorded in Section 2. The preceding first-exit argument is still needed to confine graph paths.

Use the annulus of Section 32, with its observation buffer disjoint from the pin. Lemma 32.3 identifies its normalized crossing limit as the same C times the retained reference crossing. The positive continuous density in Section 33 gives its unique median. Lemma 33.1 therefore removes C from every subsequence and gives convergence in probability on the original field–reference coupling. Together with (29.1), it gives $a(I) = I^{-\alpha+o(1)}$. Finally Lemma 33.2 gives $a(tI)/a(I) \to t^{-\alpha}$. Regular variation has not been used in the stopping or identification arguments.

Remark 29.7 (Assembly of the propagation inputs). The theorem proves the all-intensity and scalar-normalization steps once (L), (R), and the marked compactness/closure interface hold. The graph exponent and the Gaussian density calculations do not substitute for those inputs. In particular, the physical shape/coarse-history transfer in (L) and the full local-passage rigidity in (R) are supplied by the preceding dedicated sections. For the full proposed theorem their reference inputs are supplied by the cited results in Section 2 and Lemma 22.3 for every $d \geq 2$. The same two-corridor forcing lemma, rigidity theorem and propagation proof therefore cover the plane and all higher dimensions.

30 Uniform finite multiplier windows before regular variation

The successful level immediately before a first failure controls the failed level's full multiplier window through a fixed fourfold window. The required control follows from Gaussian transport, without continuity of an individual Poisson cloud or regular variation of its normalizer.

Lemma 30.1 (Gaussian response and uniform normalizer profiles). *Let $I_n \to \infty$, put $B_n = b(I_n)$ and $U_n = F_{I_n}/B_n$, and suppose that the full base experiments have a controlled joint limit retaining U_n and their fixed Gaussian background. Assume, for some $p > 0$,*

$$\mathbb{E}U_n^p = 1, \qquad \sup_n \mathbb{E}U_n^{2p} < \infty.$$

On each finite local projection, suppose the exact intensity response is the fixed Cameron–Martin change described in (29.3). Then on every compact $T \Subset (0,\infty)$ the profiles

$$r_n(t) = \frac{b(tI_n)}{b(I_n)}$$

are bounded above and away from zero, are equicontinuous, and converge uniformly to one positive continuous profile $r(t)$ along the given controlled joint subsequence.

For every deterministic $t_n \to t \in (0,\infty)$, the full local experiment at $t_n I_n$ in units B_n converges to its fixed t transported limit. If a higher base moment of order $q_* > q$ is uniformly bounded, the transported qth calibration powers are uniformly integrable on each compact T after choosing the Hölder exponent strictly close enough to one.

Proof. Fix a finite projection and one smooth cutoff equal to one on all its observation and internal-reference buffers. For $c = \gamma^{-1}\log t$, write its Gaussian likelihood as $R_c = \exp(cZ - c^2\sigma^2/2)$, where Z is a fixed centered Gaussian of variance σ^2. The exact response gives

$$r_n(e^{\gamma c})^p = \mathbb{E}[R_c U_n^p].$$

On a compact interval of c, the functions R_c are uniformly bounded in every finite L^q norm and are Lipschitz in L^2. Hence Cauchy–Schwarz gives a uniform upper bound and

$$|r_n(e^{\gamma c})^p - r_n(e^{\gamma c'})^p| \le \|R_c - R_{c'}\|_2 \|U_n^p\|_2 \le C|c - c'|.$$

Put $M = \sup_n \mathbb{E}U_n^{2p}$. Choose $\delta > 0$ so $\sup_c \mathbb{P}(R_c < \delta) < 1/(4M)$ on that interval. Then

$$\mathbb{E}[R_c U_n^p] \ge \delta(1 - \mathbb{E}[U_n^p \mathbf{1}_{\{R_c < \delta\}}]) \ge \delta/2.$$

Taking the positive pth root preserves equicontinuity on the resulting compact interval of values.

For each fixed c, the joint limit and uniform integrability give $\mathbb{E}[R_c U_n^p] \to \mathbb{E}[R_c U^p]$. To check uniform integrability directly, choose $0 < \varepsilon < 1$ and Hölder exponents with $(1+\varepsilon)pa = 2p$ and $a > 1$; the Gaussian factor has the required finite conjugate moment. Thus $\sup_n \mathbb{E}(R_c U_n^p)^{1+\varepsilon} < \infty$, uniformly for compact c. Truncating the likelihood and the calibration coordinate permits passage through joint weak convergence. Equicontinuity and a finite-net argument now upgrade pointwise profile convergence to uniform convergence on the compact interval.

For moving multipliers, the exact finite experiment is obtained by weighting the base law by R_{c_n} and then translating the field back by c_n times the cutoff. The internal reference readouts whose domains lie in its plateau acquire the scalar factor $e^{-\xi c_n}$; the retained Gaussian coordinates acquire their deterministic Cameron–Martin translations. These operations are jointly continuous in c_n on their finite projected state spaces. Any global ambient reference coordinate follows its actual compact-cutoff Weyl map, not the plateau scalar. The fixed-background measurable-recoding argument and the uniform Weyl order bound in the proof of Lemma 29.1 justify this extra coordinate, without properness of the auxiliary metric. Also $\|R_{c_n} - R_c\|_1 \to 0$, uniformly with respect to the fixed background marginal. Truncate the likelihood, use tightness of the base states and continuity on compact sets, and then remove the truncation. This proves convergence of every fixed finite projection. Finite-projection localization and a countable determining family then identify the actual whole local array. Different cutoff choices are consistent because their finite identities all come from the same actual Poisson experiment.

Finally, if $q < q_*$, choose $\varepsilon > 0$ and $a > 1$ so $q(1+\varepsilon)a < q_*$. Hölder and Gaussian moments bound $\mathbb{E}[R_c U_n^{q(1+\varepsilon)}]$ uniformly on compact c. This proves the stated uniform integrability. No density bound uniform in an unbounded multiplier range is asserted. □

Lemma 30.2 (One scalar ratio, not two)**.** *Under the preceding lemma, suppose the base limit is intrinsically $C_1 D_X$ and every fixed multiplier's own-unit limit has scalar $C_t \in [C_-, C_+]$. Then*

$$C_t r(t) = C_1 t^\alpha, \qquad R_0^{-1}(u/v)^\alpha \le \frac{r(u)}{r(v)} \le R_0 (u/v)^\alpha, \quad R_0 = C_+/C_-.$$

The profile need not yet be a power.

Proof. In the finite transported experiment the raw passage coordinates remain in base units. On a testing buffer where the cutoff is constant, the reference Weyl rule multiplies its metric by t^α. Thus the fixed-t candidate in base units has scalar $C_1 t^\alpha$ against the target reference. Dividing by the own-unit ratio $r(t)$ proves the first identity. Divide that identity at u and v; the common C_1 cancels. The factor is $C_v/C_u \in [R_0^{-1}, R_0]$, proving the claim. □

Lemma 30.3 (The failed level's whole window)**.** *Let $A_j = b(2^j)$. Suppose j_n are controlled preceding levels, their full own-unit laws converge to a rigid law, and every fixed multiplier has a retained own-unit calibration and scalar in $[C_-, C_+]$. Let $f_n = j_n + 1$. Then uniformly for $t \in [1/2, 2]$,*

$$\frac{b(t2^{f_n})}{A_{f_n}} = \frac{r_n(2t)}{r_n(2)} \longrightarrow \frac{r(2t)}{r(2)}.$$

The limiting ratio lies in $[R_0^{-1} t^\alpha, R_0 t^\alpha]$. Also, for every fixed integer $m \ge 1$,

$$\lim_n \frac{A_{f_n - m}}{A_{f_n}} = \frac{r(2^{1-m})}{r(2)} \le R_0 2^{-m\alpha}.$$

These facts require no success assumption at f_n.

Proof. Apply Lemma 30.1 at the preceding base 2^{j_n} on the fixed compact multiplier interval $[1, 4]$. Its limit is uniformly positive, so taking the displayed quotient preserves uniform convergence. Lemma 30.2 gives its bound with one factor R_0. For the block ratio, enlarge the fixed compact multiplier window to include 2^{1-m}, and apply the same pairwise identity. Every added multiplier is fixed before the discrete limit. □

Closing the first-failure contradiction. The preceding lemmas make the strictness checks explicit. Choose $R > R_0 2^\alpha$, set $s = \alpha/2$, and choose m_0 with $R_0 2^{-m_0(\alpha - s)} < 1/4$. Once the backward-array argument has identified the preceding law, the failed level's law converges to a member of the rigid class, so its open neighborhood test passes. Lemma 30.3 makes its full factor-two profile test strict, because all limiting values lie in a compact subinterval of (R^{-1}, R). The same lemma makes its block test strict, since

$$R_0 2^{-m_0 \alpha} < 2^{-m_0 s}/4.$$

Thus the alleged first failure passes all three defining tests. The window $[1, 4]$ here is fixed; no convergence uniform over an intensity range growing with n is used.

31 Confined records, escape control, and global metric convergence

This section isolates the final local-to-global implication. Its lower input concerns bounded paths already confined to deterministic compact sets. It makes no assumption about the ranges of unrestricted graph geodesics.

Lemma 31.1 (First-exit prefixes suffice for escape control)**.** *Consider embedded Delaunay graphs with unit costs $\epsilon_n \to 0$. Let $v_n(x)$ be the nucleus of the cell containing x. Assume the whole Euclidean diameters of cells meeting each fixed compact set tend uniformly to zero. Let D be a proper continuous length metric inducing the Euclidean topology and let $c > 0$. Suppose the confined lower-record property holds: whenever graph walks have contact traces in one fixed compact set, bounded normalized costs, and endpoints tending to x, y, every convergent subsequence of their cost coordinates has limit at least $cD(x, y)$.*

For each compact K and each finite cost bound M, all sufficiently large n have the following property: every graph walk beginning at $v_n(x)$ for $x \in K$ with normalized cost at most M stays in one fixed Euclidean ball. The radius can depend on D, K, c, M, but not on the walk or n.

Proof. Properness and the Euclidean topology imply

$$D(K, \partial B_R) \longrightarrow \infty.$$

Otherwise points $x_j \in K$ and $y_j \in \partial B_{R_j}$, $R_j \to \infty$, would have bounded $D(x_j, y_j)$. After taking $x_j \to x \in K$, continuity would put all y_j in one compact D-ball about x, a contradiction. Choose R so $K \Subset B_R$ and $cD(K, \partial B_R) > M + 1$.

Suppose a walk of cost at most M exits B_R. Stop at its first outside nucleus q_n, with preceding nucleus p_n inside the ball. Their cells share a facet. Since the cell of p_n meets $\overline{B}_R$, its whole diameter tends to zero. A facet point z_n therefore lies within $o(1)$ of p_n, and hence in B_{R+1} for large n. The cell of q_n also meets this fixed compact, so its diameter tends to zero. Consequently

$$|p_n - q_n| \longrightarrow 0.$$

Thus q_n lies within $o(1)$ of ∂B_R. All earlier nuclei lie in B_R, and their cells have vanishing whole diameters. The contact trace of this prefix therefore lies in the fixed compact $\overline{B}_{R+1}$. Its cost is at most M; an endpoint convention can add at most one edge, which tends to zero.

If escaping walks exist for infinitely many n, take their starting points $x_n \in K$ and first-exit endpoints. After a subsequence, $x_n \to x \in K$ and $q_n \to q \in \partial B_R$. Their bounded costs have a convergent subsequence. Confined lower capture then forces that cost to be at least $cD(x, q) > M + 1$, a contradiction. Only the stopped compact prefixes were captured. □

The first-exit argument does not require a bound for a straight Delaunay edge whose two nuclei both lie outside the ball. It uses only a cell based at the last inside vertex and the adjacent cell sharing its facet.

Lemma 31.2 (Global identification from confined upper and lower records)**.** *In the preceding setting, write $d_n(x, y)$ for the normalized unrestricted graph distance between $v_n(x), v_n(y)$. Assume also:*

1. *For each fixed pair x, y and every $\eta > 0$, confined upper records realize walks of cost at most $cD(x, y) + \eta + o(1)$.*
2. *On each compact K, a common local joining modulus bounds $d_n(x, y)$ uniformly when $x, y \in K$ have small Euclidean separation; the modulus tends to zero with that separation.*

Then $d_n \to cD$ locally uniformly.

Proof. Fix compact K and a finite sufficiently fine endpoint net. The upper records for its finitely many pairs, together with the local joining modulus, bound $\sup_{x,y \in K} d_n(x, y)$ by one finite cost

bound for all large n. The first-exit lemma now confines every shortest walk between these endpoint cells to one fixed larger compact.

If the lower convergence failed uniformly, there would be $x_n, y_n \in K$ and $\eta > 0$ with $d_n(x_n, y_n) \leq cD(x_n, y_n) - \eta$. Take a subsequence with $x_n \to x$, $y_n \to y$ and with the bounded cost coordinates convergent. The shortest walks are confined by the preceding argument. Lower capture gives a limiting cost at least $cD(x, y)$, whereas continuity of D gives an upper bound $cD(x, y) - \eta$. This is impossible.

For the upper convergence, join x, y to nearby net points, use their fixed-pair confined upper record, and join back. The error is bounded by twice the joining modulus, the continuity modulus of cD on $K \times K$, and the finitely many upper-record errors. First let n increase and then refine the net. This proves uniform upper convergence. □

Application to the random marked laws. The hypotheses of the two lemmas are consequences of the identified confined record structure, before unrestricted distances are identified. A bounded contact trace in a compact set has a captured record, and intrinsic lower comparison gives its cost bound. Realizable upper records and the adaptive endpoint-joining modulus give the upper hypotheses.

For the probability statement one need not choose a deterministic radius separately on each sample. First choose a deterministic cost bound M using tightness of the compact upper connections. The canonical proper reference and the deterministic lower bound $c \geq c_-$ then permit a deterministic R with

$$c_- D_h(K, \partial B_R) > M + 1$$

outside an event of arbitrarily small probability. At this fixed R, any failure of confinement supplies a bounded-cost first-exit prefix in one fixed enlargement. The controlled marked convergence captures that prefix, contradicting the strict displayed lower bound. This proves tightness of the unrestricted near-minimizer ranges, rather than assuming it.

One can equivalently apply the deterministic lemmas in a controlled almost-sure representation of the full marked laws. Every use of compact capture is then on a fixed buffer; no compactness of the entire original unconfined walk is invoked. The reference metric is retained jointly, so the conclusion identifies the global metric on the original field coupling after the scalar calibration.

32 GMC normalization and parameter calibration

This section proves necessary consequences of a metric convergence statement and identifies the local crossing used for normalization. Throughout, the logarithmic coefficient of the field covariance is one, $d \geq 2$, and $0 < \gamma < \sqrt{2d}$. The background-charge notation $Q_d(\xi)$ is that used for exponential metrics; existence or tightness results for those metrics do not themselves imply convergence of a Poisson–Voronoi discretization [8, 7].

32.1 The GMC convention and the target parameter

Fix a rotationally invariant regularization. We call

$$\mu_h^{\mathrm{can}}(dx) = \lim_{\varepsilon \downarrow 0} \varepsilon^{\gamma^2/2} e^{\gamma h_\varepsilon(x)}\, dx$$

the canonical chaos; an inessential fixed regularization constant is allowed. Suppose that, locally uniformly, $\operatorname{Var} h_\varepsilon(x) = \log \varepsilon^{-1} + v(x) + o(1)$. The full variance-normalized chaos satisfies

$$\mu_h^{\mathrm{var}} = e^{-\gamma^2 v/2} \mu_h^{\mathrm{can}} = \mu^{\mathrm{can}}_{h-(\gamma/2)v}. \tag{32.1}$$

Thus the two conventions are not interchangeable without a deterministic field correction. The usual subcritical construction and its covariance comparison properties apply here [9]. In the rest of this section μ_h denotes the canonical convention.

For a similarity $\phi(x) = rOx + z$, where $r > 0$ and O is orthogonal, a change of variables at regularization scale $r\varepsilon$ gives

$$\mu_{h\circ\phi}(A) = r^{-d-\gamma^2/2} \mu_h(\phi(A)), \qquad \mu_{h+f} = e^{\gamma f} \mu_h. \tag{32.2}$$

In contrast, the pullback exponent of μ^{var} is d; indeed $v_{h\circ\phi} = v_h \circ \phi - \log r$. Equation (32.1) reconciles these statements. Set

$$q_\mu = \frac{d}{\gamma} + \frac{\gamma}{2}, \qquad h^\phi = h \circ \phi + q_\mu \log r. \tag{32.3}$$

Then $\mu_{h^\phi}(A) = \mu_h(\phi(A))$. Consequently the pullback of a Poisson cloud of intensity $\lambda\mu_h$ has intensity $\lambda\mu_{h^\phi}$, and its Euclidean Voronoi diagram and Delaunay graph are the corresponding pullbacks. There is no intensity rescaling in this last statement.

There is also an exact constant-shift identity. Conditional on h, take a Poisson random measure on $\mathbb{R}^d \times (0,\infty)$ with intensity $\mu_h(dx)\, du$, and retain (x,u) when $u \leq \lambda$. If $f = c$ on an open set W, the clouds for $h + f$ at intensity λ and for h at intensity $e^{\gamma c}\lambda$ can be coupled to coincide on W:

$$\mathcal{P}_\lambda(h+f) \cap W = \mathcal{P}_{e^{\gamma c}\lambda}(h) \cap W. \tag{32.4}$$

This identity does not assert monotonicity of Delaunay graph distances.

Proposition 32.1 (Necessary background charge)**.** *Suppose the graph metrics converge locally uniformly, jointly with the field and with the same deterministic normalization, to a positive continuous local length metric D_h inducing the Euclidean topology. Suppose the target has smooth Weyl coefficient $\xi > 0$ and similarity rule*

$$D_{h\circ\phi + Q\log r}(x,y) = D_h(\phi(x), \phi(y)).$$

Assume the convergence also transfers to h^ϕ on local buffers, either as a hypothesis or by local absolute continuity of the field laws. Then necessarily

$$Q = q_\mu = \frac{d}{\gamma} + \frac{\gamma}{2}. \tag{32.5}$$

In particular a limit with background charge $Q_d(\xi)$ must satisfy $Q_d(\xi) = d/\gamma + \gamma/2$.

Proof. Similarity covariance of the discrete diagram, followed by convergence on a fixed annulus, gives $S_A(D_{h^\phi}) = S_{\phi(A)}(D_h)$, where S_A is its positive finite crossing distance. The target similarity rule and constant Weyl scaling give instead

$$S_A(D_{h^\phi}) = r^{\xi(q_\mu - Q)} S_{\phi(A)}(D_h).$$

Choose $r \neq 1$ and cancel the positive crossing distance. Localized annular crossings suffice; their localization is proved in Lemma 32.3 below. □

Remark 32.2 (Pinning and scope). Local absolute continuity in Proposition 32.1 cannot be replaced by an assertion that differently pinned whole-space fields have equivalent global laws. A sufficient local argument is as follows. Take a buffer disjoint from both pinning supports, and assume smooth cutoff functions equal to one on the buffer are Cameron–Martin directions for each pinning. Adding an independent nondegenerate Gaussian constant to either pinned version produces equivalent unpinned laws: conditional on the common field modulo constants, both constants have everywhere positive Gaussian densities. On the buffer, adding any fixed constant has an equivalent law by the Cameron–Martin cutoff, so mixing those shifts preserves equivalence in both directions. The two pinned restrictions are therefore equivalent. This applies to the usual whole-space logarithmic field and its similarities on such buffers.

Equation (32.5) is a necessary calibration, not a proof of graph convergence or of existence of a solution ξ for every proposed γ. For $Q > \sqrt{2d}$, the subcritical solution for the chaos parameter is $\gamma = Q - \sqrt{Q^2 - 2d}$.

32.2 Local crossings and normalization

Let $d_\lambda^h(x, y)$ be the unit-edge Delaunay distance between the selected cells containing x, y, with the fixed measurable boundary convention. Work with locally finite clouds in general position. Fix an annulus $A = \{r_- < |x - z| < r_+\}$ with compact closure in an open buffer W. A cell-trace crossing is a continuous curve across A, perturbed to pass through facet interiors and avoid lower-dimensional faces. Let $C_\lambda^h(A)$ be the least number of cell changes along such a crossing. Endpoint conventions may change this by a fixed constant. Put

$$G_\lambda^h(A) = \inf_{\substack{x \in \partial B_{r_-}(z) \\ y \in \partial B_{r_+}(z)}} d_\lambda^h(x, y).$$

Lemma 32.3 (Crossing comparison and localization)**.** *With consistent endpoint conventions, $|C_\lambda^h(A) - G_\lambda^h(A)| \leq 2$. If μ_h is almost surely locally finite and has full support, then, with probability tending to one, $C_\lambda^h(A)$ is determined solely by $\mathcal{P}_\lambda(h) \cap W$. The same conclusion holds under any fixed smooth field shift.*

Proof. A generic crossing lists a facet-adjacent cell chain, so it bounds the graph set distance from above. Conversely realize a graph chain by a curve: join successive facet-interior points inside the intervening convex Voronoi cells and choose its endpoints on the two spheres. Trim between its last visit to the inner sphere before its first subsequent visit to the outer sphere. The resulting curve stays in the closed annulus and visits no additional cells. Generic perturbations inside the selected cells are possible. At each endpoint the selected cell and any incident cell used by a perturbation are equal or adjacent: the general-position property makes the incident cells a Delaunay simplex. Thus one additional edge at each endpoint suffices, giving the bound two. With generic endpoint conventions from the outset the two minima agree.

Choose a fixed mesh size $\eta > 0$ small compared with $\operatorname{dist}(\overline{A}, W^c)$. Cover a neighborhood of $\overline{A}$ by finitely many balls of radius η, all lying in W. Each ball has positive μ_h-mass almost surely, and its probability of being empty is $\mathbb{E}[e^{-\lambda\mu_h(B)}] \to 0$. On the filling event every point near $\overline{A}$ is closer to a nucleus in W than to any nucleus outside W. The diagram restricted to that neighborhood, hence its cell-trace crossing, is determined by the restricted cloud. Multiplication of the measure by a fixed positive smooth density preserves the argument. □

The scalar calibration and intensity ratio are proved in Section 33, after subsequential identification. Neither normalization limit is an input to the graph compactness or propagation arguments.

33 Normalizer calibration and convergence in probability

The arguments here begin after the subsequential metric identification in Sections 28–31. They do not use regular variation to obtain that identification. In particular, the scalar under the moment normalization need not be the same along different subsequences.

Lemma 33.1 (One annular quantile removes every subsequential scalar)**.** *Let D be the retained reference metric, $b(I) \to \infty$ be deterministic, and C_I be the graph crossing count of one fixed buffered annulus. Write S for the reference crossing. Suppose $0 < S < \infty$ almost surely and its distribution function is continuous and strictly increasing on $(0, \infty)$. Assume every deterministic sequence $I_n \to \infty$ has a subsequence for which, jointly with the retained reference,*

$$(b(I_n)^{-1} d_{I_n},\ b(I_n)^{-1} C_{I_n},\ D) \Longrightarrow (cD, cS, D),$$

locally uniformly in the metric coordinate, for some deterministic $c \in [c_-, c_+] \Subset (0, \infty)$. Let m_I, m_S be the lower medians of C_I, S. For all sufficiently large I, $0 < m_I < \infty$, and $a(I) = m_S/m_I$ satisfies

$$a(I) d_I \longrightarrow D \quad \textit{locally uniformly in probability on the original coupling.}$$

Moreover $a(I)b(I)$ stays bounded above and away from zero. Therefore $b(I) = I^{\alpha+o(1)}$ implies $a(I) = I^{-\alpha+o(1)}$.

Proof. Along one identified subsequence, ordinary quantile convergence gives $m_{I_n}/b(I_n) \to cm_S$. Indeed, for every small $\varepsilon > 0$, continuity and strict increase put the distribution function of cS strictly below $1/2$ at $cm_S - \varepsilon$ and strictly above $1/2$ at $cm_S + \varepsilon$. Weak convergence gives the same two strict inequalities eventually for the approximating distribution functions. They bracket the lower median. No continuity of the prelimit crossing law is required.

Consequently $a(I_n)b(I_n) \to c^{-1}$, and the normalized joint limit is (D, D) in the candidate and retained reference coordinates. The local-uniform discrepancy between those two coordinates converges in distribution to zero, hence in probability to zero. Every original sequence has such a further subsequence. A deterministic sequence along which a fixed discrepancy had probability bounded below would contradict this conclusion. This proves full convergence in probability on the original coupling.

The same subsequence argument shows that $m_I/b(I)$ eventually lies in a fixed compact subinterval of $(0, \infty)$: otherwise a sequence violating such a bound would have a subsequential limit cm_S with $c \in [c_-, c_+]$. This proves eventual positivity and finiteness of the graph median as well. Arbitrary positive definitions at the bounded set of earlier intensities do not affect the conclusion. The growth exponent follows. □

33.1 Why the annular median has no flat quantile

Fix the pin used for the actual graph experiment. Choose the annulus and its whole observation buffer W disjoint from the pin support. A smooth Cameron–Martin direction ψ can then equal one near W and vanish on the pin support. Let $S_A(D)$ denote the infimum of reference lengths of annular crossings, with endpoints on its two boundary spheres.

This observable equals the unrestricted metric set distance between the two spheres. Given any connecting path, cut between its last inner-sphere visit before its first outer-sphere visit. The retained path stays in the closed annulus and in its open interior away from the two endpoints. Conversely every annular crossing is an unrestricted connection. Compactness of the disjoint spheres and continuity of the positive metric imply $0 < S_A(D) < \infty$.

Although unrestricted point distances need not be local, this particular set distance depends only on the internal metric near the annulus. Simultaneous Weyl covariance therefore gives

$$S_A(D_{h+t\psi}) = e^{\xi t} S_A(D_h).$$

Gaussian projection in the Cameron–Martin direction writes $h = h^{\perp} + Z\psi$, where Z is an independent nondegenerate Gaussian. The scalar law consequently has the form

$$S_A(D_h) = e^{\xi Z} T, \qquad 0 < T < \infty,$$

where T is measurable with respect to $h^{\perp}$. No identification of globally differently pinned laws is used.

If the metric has only been specified under the original Gaussian law, this assertion still follows without evaluating it under the singular law of $h^{\perp}$. Intersect the Weyl identities for rational shifts and disintegrate with respect to $h^{\perp}$. The function $t \mapsto \log S_A(D_{h^{\perp}+t\psi}) - \xi t$ is almost everywhere in t invariant under every rational translation. A measurable such function is almost everywhere constant. To see this, first apply a bounded injective transform and then use continuity of translations in local L^1 to extend invariance from rational to real shifts. That constant defines $\log T$ measurably.

The law of $\log S_A(D_h)$ is the convolution of the law of $\log T$ with a nondegenerate Gaussian. Its density is continuous and strictly positive everywhere. Hence the law of $S_A(D_h)$ has continuous strictly increasing distribution function on $(0, \infty)$, and every interior quantile is unique.

33.2 The graph normalizer ratio after full probability convergence

Lemma 33.2 (Intensity response determines the normalizer ratio)**.** *Suppose the already identified graph metrics satisfy $a(I)d_I \to D_h$ locally uniformly in probability, with $a(I) \to 0$, using canonical GMC and the preceding pin-compatible annular crossing. Then, for every fixed $t > 0$,*

$$\frac{a(tI)}{a(I)} \longrightarrow t^{-\xi/\gamma}.$$

Proof. Write $t = e^{\gamma c}$. The fixed Cameron–Martin translation $h \mapsto h + c\psi$ has a fixed integrable likelihood, so events whose probabilities tend to zero before this translation still have probabilities tending to zero afterward. This applies to the joint field and its conditional Poisson kernel. The established convergence under that likelihood, together with Lemma 32.3 and $a(I) \to 0$, gives

$$a(I)C_I^{h+c\psi}(A) \longrightarrow e^{\xi c}S_A(D_h)$$

in probability on the original shifted experiment.

Canonical GMC obeys $\mu_{h+c\psi} = e^{\gamma c}\mu_h$ on the whole buffer. Couple the two Poisson restrictions there to be identical. The finite-buffer filling event determines the cell-trace crossing from that restricted cloud with probability tending to one. Therefore

$$a(I)C_{tI}^h(A) \longrightarrow e^{\xi c}S_A(D_h), \qquad a(tI)C_{tI}^h(A) \longrightarrow S_A(D_h).$$

These convergences hold jointly on the same field coupling. Their positive common crossing permits division. The quotient is the deterministic number $a(I)/a(tI)$ on an event whose probability tends to one, so this number tends to $e^{\xi c} = t^{\xi/\gamma}$. □

The lower graph median is a Borel function of intensity: for each integer k, the probability that the measurable Poisson crossing count is at most k is Borel in I. Taking its lower $1/2$ quantile preserves Borel measurability. Thus the measurable regular-variation theory applies. The power exponent has already been obtained from the moment normalization and the compact scalar interval.

33.3 The internal probability-convergence statement

Proposition 33.3 (Convergence in probability and regular variation)**.** *Under the reference-metric inputs of Section 2, let $d \geq 2$ and $0 < \gamma < \sqrt{2d}$. With $a(I)$ defined by the annular median in (1.4), for every compact $K \subset \mathbb{R}^d$,*

$$\sup_{x,y\in K} |a(I)d_I(x,y) - D_h(x,y)| \xrightarrow[I\to\infty]{\mathbb{P}} 0. \tag{33.1}$$

The function a is positive and Borel measurable, and satisfies

$$a(I) = I^{-\alpha+o(1)}, \qquad \frac{a(tI)}{a(I)} \longrightarrow t^{-\alpha} \quad (t > 0), \qquad \alpha = \frac{\xi}{\gamma}.$$

Proof. Section 2 supplies the reference metric and its quantitative inputs. Lemma 22.3 verifies the common anchor property, and Lemma 27.2 verifies the two-corridor construction, in both the planar and higher-dimensional cases.

The screening, joining, cell-size, and Poisson-count estimates give the moment normalization $b(I) = I^{\alpha+o(1)}$, vanishing meshes, and the marked compactness interface. The common local kernels and Gaussian transfer provide the conditional-law hypotheses of Theorem 28.4. Moment normalization and uniform integrability retain a nonzero guarded port calibration. The annular joining, comparison, and forcing arguments establish the remaining hypotheses, giving restricted rigidity for every $d \geq 2$.

The physical-label transfer and this rigidity are inputs (L) and (R) of Theorem 29.3. That theorem gives subsequential identification in the original deterministic units, with deterministic scalars in one compact subinterval of $(0, \infty)$. The first-exit and local-to-global lemmas of Section 31 identify unrestricted distances jointly with the actual field and reference metric. The crossing observable is retained in this joint identification: by Lemma 32.3 it differs by at most two hops from the metric set distance between two fixed compact spheres, whose locally uniform limit is their reference set distance.

The annular crossing law has a unique positive finite median, as proved above. Lemma 33.1 therefore gives the normalization (1.4) and (33.1) on the original coupling. Its comparison with $b(I)$ gives the power exponent and in particular $a(I) \to 0$. Only after establishing this probability convergence do we apply Lemma 33.2, which gives the ratio limit. The preceding measurability argument completes the proof. □

34 Local tests and almost sure small-distance control

Previously established conclusions used below. Write $D = D_h$, $\mu = \mu_h^\gamma$, $X_I = a(I)d_I$, $\alpha = \xi/\gamma$ and $\Delta = 1/\alpha$. Proposition 33.3 has established $X_I \to D$ locally uniformly in probability on the original coupling, with the annular-median normalization (1.4) and the regular-variation estimates (1.6). The function a is positive and measurable. Section 2 supplies the proper continuous reference length metric, its local Euclidean upper Hölder bound, its local internal-metric readouts, and its Weyl and similarity identities. The upper metric-volume estimate in (2.2), also recorded in Proposition 14.1, holds for every exponent loss $\eta > 0$, almost surely eventually as $r \downarrow 0$, uniformly for centers in any fixed compact buffer.

The whole-cell estimate (15.4) and the Euclidean localization in its proof hold simultaneously over the marked process. They also hold on arbitrary slabs $[s, 2s]$: choose $2^n \le s < 2^{n+1}$ and use the two adjacent dyadic slabs. Applying their metric-cell bound with loss $\kappa/2$ and absorbing the fixed factors gives whole D-diameter at most $s^{-\alpha+\kappa}$ for every cell meeting a fixed compact buffer, uniformly throughout $[s, 2s]$, for all sufficiently large s. Their whole Euclidean diameters also tend uniformly to zero. Lemmas 12.2 and 14.3 give local filling and facet connectivity of the complete trace-cell family; Lemma 12.4 bounds the number of cells incident to a point by $d + 1$. Common local reference rules and the local Gaussian transfers are established in Sections 16, 18 and 21. Finally, Lemma 24.1 supplies the joint Gaussian comparison of annular shapes and coarse histories; Section 35 below proves its finite-intensity application, including a smooth pin that meets the compact under consideration.

The remaining sections use these internal conclusions to prove the almost sure assertion. The annular improvement and normalizer-selection strategy is related to Devlin [16]; the finite-intensity quantum labels and the interpolation of changing Delaunay graphs are proved below.

For an open set U, a confined graph route is required to have all its cells in U. Fixed inner and outer buffers make this equivalent for the purposes below to confining its nuclei or embedded edges. Put $A_{u,v}(z) = B_v(z) \setminus \overline{B}_u(z)$.

Lemma 34.1 (Buffered local comparisons)**.** *Fix compact annular regions strictly inside a connected open annular observation window. On these regions, qualitative locally uniform graph convergence and vanishing cell mesh have the following consequences. For every $\varepsilon > 0$, with probability tending to one as $J \to \infty$, each reference path crossing two fixed separated boundary strips, and confined to the compact middle region, has a graph follow in the larger window of $a(J)$-cost at most $(1 + \varepsilon)$ times its reference length. Conversely, the endpoints of each graph crossing confined to the middle region can be joined inside the larger window at reference cost at most $(1 + \varepsilon)$ times its $a(J)$-cost.*

The ratios of the intrinsic diameter of a fixed inner boundary strip to the intrinsic distance between the two strips are tight, for both the reference and the normalized graph metric. Their high-probability bounds can be chosen independently of ε. For any deterministic $u_J \downarrow 0$, the normalized local graph diameter of pairs at Euclidean distance at most u_J tends to zero in probability, uniformly over the fixed strips. More precisely, if $\omega_J(t)$ and $\omega_D(t)$ are the two local moduli at Euclidean radius t, and q_J, q_D are their positive strip-crossing costs, then for every $\varepsilon > 0$,

$$\lim_{t \downarrow 0} \limsup_{J \to \infty} \mathbb{P}[\omega_J(t) > \varepsilon q_D \text{ or } \omega_D(t) > \varepsilon q_J] = 0. \tag{34.1}$$

These assertions can be expressed by events depending only on the field and Poisson cloud in the observation annulus.

Proof. We give the deterministic localization argument behind these assertions. Choose nested

compact buffers around the middle region, still inside the observation window. On an event of arbitrarily high probability the reference distance between the two strips is bounded below by $c > 0$, the reference distance from the inner buffers to the exterior of the larger window is bounded below by $c' > 0$, and their relevant intrinsic diameters are bounded above by a finite deterministic cutoff. Positivity, continuity and the length property justify these cutoffs.

Take $v < c'/10$. If the uniform discrepancy between the ambient normalized graph distance and D on the surrounding compact is e, a graph geodesic of cost at most $2v$ starting in the inner buffer cannot leave the larger window when e and the cell mesh are small. Its first-exit prefix would give a pair with reference distance at least $c' - o(1)$ and graph distance at most $2v$. Reference paths of length at most $2v$ satisfy the same confinement assertion directly.

Partition a source path of length L into pieces of source cost at most v. For graph paths use vertex cuts; the final overshoot is at most one edge cost, which tends to zero. Replace each piece by a short target path. The preceding exit argument confines the replacements to the larger window. There are at most $2L/v + 1$ pieces for all large J, and their total cost is at most

$$L + e(2L/v + 1) + o(1)(L/v + 1).$$

Every crossing under consideration has cost at least $c/2$ for large J. The displayed error is therefore at most εL once the discrepancy and edge cost are sufficiently small. This works simultaneously for all the specified paths. In the reverse direction the source graph trace is kept in a fixed compact middle annulus; no continuity of internal metrics at the boundary of an open observation set is assumed.

For the diameter bound in a specified open annulus, choose a finite reference connection network compactly contained in that annulus, together with finitely many short neighborhoods covering the strip. This is possible by compactness, local continuity and connectedness of the annulus. Following this network, with its compact clearance from the boundary, bounds all the graph joins. On events of arbitrarily high probability its lengths and inverse clearances have fixed cutoffs. The crossing lower bound follows by the reverse construction. The reference diameter is finite and its crossing distance is positive, so their ratio has a finite deterministic cutoff at any fixed probability level. This cutoff does not depend on the later comparison accuracy. Small-distance control follows by first choosing the reference modulus and then taking the uniform discrepancy small, with the same exit argument. The reference crossing is positive and the graph crossing is tight away from zero, so this also proves the two-parameter formulation (34.1).

Define the events using intrinsic target costs in the larger window and source paths in the compact middle region. Local reference readouts and the local filling event make them measurable from that annulus alone. Ambient convergence has been used only to prove their probability tends to one. To obtain the unit sphere-pinned annular law used below, first apply this argument in one deterministic translate of the observation annulus disjoint from the original pin support. Pull back by translation; this changes neither the intensity nor $a(J)$. The translated field law on the unit annulus is equivalent to the radius-eight sphere-pinned law: both pins are exterior, so Lemma 18.1 compares each to the same local finite-range Gaussian law. Append the same local Poisson kernel and common local reference readouts. A fixed integrable density transfers vanishing failure probabilities, with the same numerical normalizer. This transfer is used qualitatively only; it supplies no rate and makes no substitution of a random intensity. □

Lemma 34.2 (An almost sure microscopic modulus)**.** *Fix a compact inner buffer $K \Subset U$, with U open. For every $\nu > 0$, almost surely, for all sufficiently large real I, every reference- rectifiable*

path P contained in K has a graph follow in U satisfying

$$X_I(P(0),P(1);U) \le I^{\nu}\,\mathrm{len}_D(P) + I^{-\alpha+\nu}. \tag{34.2}$$

Consequently, for every $\eta>0$ and every compact K, almost surely

$$\sup_{x,y\in K,\ |x-y|\le I^{-\eta}} X_I(x,y) \longrightarrow 0 \quad \textit{through all real } I. \tag{34.3}$$

The same holds with confinement to a fixed slightly larger buffer.

Proof. Work on dyadic slabs $s=2^n\le I\le 2s$. For small $\kappa,\tau>0$ put $r=s^{-\alpha+\kappa}$. Conditional on h, the upper cloud $\mathcal{P}_{2s}$ satisfies, almost surely eventually,

$$\sup_{z\in K'} \#(\mathcal{P}_{2s}\cap B_D(z,4r)) \le s^{\kappa\Delta+\tau} \tag{34.4}$$

on each fixed enlarged buffer K'. To verify this, take a deterministic Euclidean grid of mesh s^{-M}, with $M\sigma>\alpha-\kappa$. The local Hölder bound makes it a D-net at scale r, eventually for almost every field. By (2.2), the conditional Poisson means of the required grid-centered balls are at most $s^{\kappa\Delta+\tau/2}$ for all sufficiently large s. The Poisson upper tail at $s^{\kappa\Delta+\tau}$, followed by the polynomial-size grid union bound, is summable. Conditional Borel–Cantelli proves (34.4). The field-dependent onset causes no difficulty; only the remaining tail of each conditional series is used.

The slab cell estimate with loss $\kappa/2$ makes every cell meeting P have D-diameter less than r. Divide P into at most $1+\mathrm{len}_D(P)/r$ pieces of D-length at most r. All nuclei of cells meeting one piece lie in a D-ball of radius $2r$. Their complete trace-cell family is facet connected, so (34.4) gives

$$d_I(P(0),P(1);U) \le C\big(1+\mathrm{len}_D(P)s^{\alpha-\kappa}\big)s^{\kappa\Delta+\tau}.$$

Whole Euclidean cell mesh tends uniformly to zero, which supplies the confinement to U. The estimate is simultaneous over all these paths and all I in the slab. Multiply by $a(I)=I^{-\alpha+o(1)}$ and choose κ,τ so that both $\kappa\Delta+\tau$ and $\kappa(\Delta-1)+\tau$ are less than $\nu/2$. The exponent slack absorbs the fixed factors and proves (34.2).

For close points of a fixed compact, short reference paths remain in a fixed enlargement: the reference clearance from the compact to its exterior is positive. Such paths can have length arbitrarily close to $D(x,y)\le C|x-y|^{\sigma}$. In (34.2), choose $\nu<\min\{\alpha/2,\eta\sigma/2\}$. This proves (34.3). All assertions are placed on one event by using rational losses and integer-radius buffers. □

35 Quantitative annular coverage at finite intensity

Fix the unit observation annulus $S=A_{1/16,4}(0)$ and let σ_8 be uniform probability on $\partial B_8(0)$. For a deterministic center and radius write

$$H(z,r)=(h,\sigma_{8,z,r}), \qquad G_{z,r}(u)=h(z+ru)-H(z,r), \quad u\in S.$$

All $G_{z,r}$ have the same annular Gaussian law n. A local test is a measurable function of this field restriction, its canonical chaos, the Poisson cloud in S, and the fixed local reference readouts. Local filling is included whenever the readout is intended to agree with a whole-space graph readout.

Lemma 35.1 (Finite-intensity all-point coverage)**.** *There are constants*

$$0 < u < v < \infty, \quad c, C > 0, \quad \varepsilon_* > 0,$$

and fixed meshes and geometric radii, with the following property. Let $N \to \infty$, $I > 0$, $J_i = I2^{-i}$, and $T \subset \{N, \dots, 2N\}$ be deterministic with $|T| \geq N/4$. For each $i \in T$ suppose a local test at intensity J_i, under n, fails with probability at most ε_*. Except on an event of probability $C_K e^{-cN}$, every point of a fixed compact K lies in the central hole of a ball $B(z, r)$ with

$$e^{-vN} \leq r \leq e^{-uN}, \tag{35.1}$$

up to fixed endpoint factors, for which the corresponding test succeeds on the actual field and cloud, in the coordinates $z + r\cdot$, at a label $i \in T$. The field used in this test is

$$g = h(z + r\cdot) + Q \log r + \gamma^{-1} i \log 2. \tag{35.2}$$

The constants and the allowed unit failure probability can be fixed before choosing the accuracy of the local metric test. They apply directly to the original smooth pin, even when its support meets K.

Proof. This is the finite-resolution form of the shape/history calculation in Section 24. We give the probability and conditioning details because a limiting-law statement alone would not suffice here.

Choose a small mesh constant $c_0 > 0$ and a large fixed logarithmic spacing L. At radii $r_j = e^{-t_j}$, $t_j = t_0 + jL$, use meshes of spacing at most $c_0 r_j$. The scales range from $t_0 = uN + O(L)$ to $t_M = vN + O(L)$. Each point of the finest mesh has an ancestry of nearest coarser centers z_j, all within $c_0 r_j$ of it. Take L large enough that the annuli $z_j + r_j S$ along one ancestry are disjoint and each finer normalization sphere lies in the central hole of every preceding annulus. There are at most $C_K e^{dvN}$ ancestries. Every point of K belongs to all the central holes along an ancestry after a fixed reduction of the mesh spacing.

For one ancestry put $H_j = H(z_j, r_j)$, $G_j = G_{z_j, r_j}$ and $Z_j = H_j - H_{j-1}$. The Gaussian covariance estimate gives

$$\mathbb{E}_* \mathcal{R}^2 \leq \exp(CM/L + C_K), \tag{35.3}$$

where $\mathcal{R}$ is the density of the actual joint law of $(G_0, \dots, G_M, H_0, Z_1, \dots, Z_M)$ relative to the product of its individual Gaussian marginals. In particular, that product law makes the shapes independent of one another and of the full coarse history.

Here the restriction on the pin can be removed without an absolute-continuity loss in the rate. If the smooth pin is ρ, its covariance is

$$K_\rho(x, y) = -\log|x - y| + F_\rho(x) + F_\rho(y) - C_\rho, \qquad F_\rho = \log|\cdot| * \rho.$$

The function F_ρ is smooth everywhere, since it is the convolution of a distribution with a compactly supported smooth function. Every G_j and Z_j is a zero-mass test, so the pin terms cancel exactly in their mutual covariances. Pairings with H_0 acquire only $O_K(r_j)$ or $O_K(r_{j-1})$ corrections. Also $\operatorname{Var} H_0 = t_0 + O_K(1)$ and $\operatorname{Var} Z_j = L + O(1)$. Thus the same normalized covariance operator $\mathcal{C}$ as in Lemma 24.1 satisfies

$$\|\mathcal{C} - \mathrm{Id}\| \leq CL^{-1/2} + C_K N^{-1/2}, \qquad \|\mathcal{C} - \mathrm{Id}\|_{\mathrm{HS}}^2 \leq CM/L + C_K/N.$$

The first norm is below $1/2$ for large L, N. For eigenvalues $1 + \delta_k$, the squared-density integral is $\prod_k (1 - \delta_k^2)^{-1/2}$. Finite projections and their bounded L^2 density martingale prove (35.3). This

argument uses the smooth pin itself, so no qualitative transfer of a small probability is being used as a polynomial bound.

Set $\Lambda_j = \gamma(Qt_j - H_j)$. Choose a fixed $\zeta \in (\sqrt{2d}, Q)$, then $u > 0$ small and v large enough that

$$\gamma(Q+\zeta)u < \tfrac{1}{2}\log 2, \qquad \gamma(Q-\zeta)v > 3\log 2.$$

Gaussian union bounds, counting endpoint centers at their own scale, give $|H_0| \le \zeta t_0$ and $|H_M| \le \zeta t_M$ simultaneously for all endpoint centers, outside probability $C_K e^{-cN}$. Hence $\Lambda_0 < N\log 2$ and $\Lambda_M > 2N\log 2$ for every ancestry.

Assign each target $i\log 2$, $i \in T$, to its first upward crossing by Λ_j. This does not require monotonicity. A jump of size $s > B_0$, where $B_0 \ge 2\log 2$, receives at most $2s/\log 2$ labels. The number of labels assigned to such jumps is therefore at most

$$\frac{2}{B_0 \log 2} \sum_{j=1}^{M} (\Lambda_j - \Lambda_{j-1})^2.$$

The increment covariance has norm at most CL and trace at most CML, so $\mathbb{E}\exp(cL^{-1}\sum_j Z_j^2) \le e^{CM}$ for small fixed c. Since $\Lambda_j - \Lambda_{j-1} = \gamma QL - \gamma Z_j$, choosing B_0 large makes the probability of losing $N/8$ labels at most e^{-qN}, for any prescribed fixed q. Take $q > dv + 3$ and union over ancestries.

At most $J_0 = \lceil B_0/\log 2\rceil + 2$ retained labels use one annulus. Keeping one label per annulus leaves at least $c_1 N$ trials along each ancestry. Their assignments depend only on the coarse history. For each,

$$g = G_j + c_j, \qquad c_j = H_j - Qt_j + \gamma^{-1} i\log 2 \in [-B_0/\gamma, 0].$$

A fixed smooth Cameron–Martin shift equal to c on S and zero near the radius-eight pin has energy at most Cc^2. Consequently a unit test of failure probability ε_* has failure probability at most $C_{B_0}\varepsilon_*^{1/2}$ under any of these shifted shape laws. Appending the same conditional Poisson kernel leaves this likelihood unchanged.

These are exactly the kernels of the original cloud. Indeed canonical GMC scaling gives

$$\mu_h(z + r\,du) = 2^{-i}\mu_g(du), \qquad I\mu_h(z + r\,du) = J_i\mu_g(du). \tag{35.4}$$

Conditional on the shapes and coarse history, the actual Poisson restrictions in the disjoint physical annuli have the product of these kernels. That same product is appended to the product Gaussian law. The likelihood is still $\mathcal{R}$. Changing the assigned label in an annulus gives the same physical intensity in (35.4).

Under the product Gaussian law, condition on a coarse history satisfying the endpoint and overshoot conditions above. The selected trials are independent, and each fails with probability at most $p = C_{B_0}\varepsilon_*^{1/2}$. Thus the product probability of a good coarse history and fewer than $c_1N/2$ successes is at most $2^{M+1}p^{c_1N/2}$. Choose ε_* so small that Cauchy–Schwarz with (35.3) makes the actual probability at most $e^{-q'N}$, with $q' > dv + 3$. Union over ancestries and add the coarse endpoint and overshoot errors. This proves the stated exponential bound and, in fact, a positive multiple of N successful annuli at every point. All geometric constants and the failure threshold have been fixed before the accuracy of any metric test. □

The covariance and local-reference identities above concern fixed deterministic charts at any one finite I, N. All subsequent almost sure uses involve only countably many intensity grids and chart libraries.

36 From qualitative convergence to comparisons with constant near one

Define

$$\ell(I) = I^{\alpha} a(I), \qquad J_i = I2^{-i}, \qquad R_I(i) = \frac{\ell(I)}{\ell(J_i)}. \tag{36.1}$$

In all ball-to-ball comparisons below, the endpoint perturbation balls are closed Euclidean balls, denoted $\overline{B}_r(x)$. At a physical annulus with label i and field (35.2), similarity covariance gives $D_h(z+r\cdot, z+r\cdot) = 2^{-\alpha i} D_g$ for the internal readouts. If $w(g)$ is a reference cost in unit coordinates, its physical value is $w_B = 2^{-\alpha i} w(g)$, whereas a normalized graph cost corresponding to $a(J_i) d_{J_i}$ is multiplied by $a(I)/a(J_i)$. Thus its factor relative to the reference is precisely $R_I(i)$. No estimate of this ratio across polynomial intensity ranges is inferred merely from regular variation.

36.1 A coarse comparison and bounded normalizer ratios

Lemma 36.1 (Charged balls)**.** *Suppose a compact path range is covered by the central holes of finitely many Euclidean balls of radii at most r_*. Give each ball a weight $w_B > 0$. Suppose a completed traversal from its hole to its boundary incurs source cost at least $m w_B$, with these charges disjoint in path order, and any pair of points on its boundary has a target join of cost at most $M w_B$ in a fixed enlargement of the ball. Then the endpoints, perturbed by at most $4r_*$, have a target connection of cost at most M/m times the source path cost. Its range is within $C r_*$ of the source range.*

Proof. Greedily choose a hole containing the current point, run to the selected ball's first subsequent boundary hit, and repeat, leaving a final incomplete traversal if necessary. Only completed traversals are charged. Replace each completed ball by a maximal member of the completed list containing it. Consecutive replacements intersect. Retain a simple chain in their intersection graph. Distinct maximal balls do not contain one another; their boundaries intersect because $d \geq 2$. Join along these boundaries, using each retained ball once. Its weight was charged at least once, so the cost is at most $M \sum_B w_B \leq (M/m)$ times the source cost. The first and last boundary points give the stated endpoint error. If there is no completed traversal the endpoints lie in one ball and the two perturbed endpoint sets intersect. Every retained ball was actually traversed, which gives the range bound. This is a finite construction and does not compare the weights of nested balls. □

Choose a fixed buffered annular crossing, say from radius $1/2$ to radius 1, and let $w(g) > 0$ be its reference internal cost in a larger annulus. The internal reference diameter of the radius-one sphere is finite for $d \geq 2$. Lemma 34.1 therefore permits a fixed $H \geq 1$, chosen at the probability threshold of Lemma 35.1, with the following unit local tests for all sufficiently large J:

(a) Reference joins between all radius-one boundary points cost at most $Hw(g)$ in the buffered annulus.

(b) Graph joins between those points, for all relevant endpoint cells, cost at most $2Hw(g)$ in $a(J)$ units.

(c) Every graph crossing whose first and last cells meet the radius-$1/2$ and radius-one spheres, and whose clipped cell trace stays in a fixed compact enlargement of that shell, costs at least $w(g)/2$ in $a(J)$ units. The compact enlargement lies strictly inside the observation

annulus. Also $C_d a(J) \le w(g)/100$, for a fixed sufficiently large dimensional constant C_d. The general-position input and facet connectivity let this fixed constant absorb every change of endpoint cell at a tie.

(d) Filling makes these readouts exactly local and prevents a single cell or edge from jumping the tested shell.

All of these can be imposed on an event of failure probability at most ε_*. The reference and graph joining bounds may be enlarged by harmless fixed factors in choosing H.

For a graph walk of $n \ge 1$ edges, choose a continuous trace through its successive convex cells, crossing the specified facets. Give its first and last cell portions time $1/2$ each and its interior cell portions time 1 each, for total time n. A subtrace crossing the shell has parameter length at least its clipped graph crossing cost in hops minus a fixed dimensional number of hops. Condition (c) absorbs this loss. Hence the local tests imply that every completed graph traversal has a disjoint charge of at least $R_I(i)w_B/4$ in X_I units. Reference traversals have charge at least w_B. Applying Lemma 36.1 on the coverage event gives, with $\rho_N = Ce^{-uN}$,

$$X_I(\overline{B}_{\rho_N}(x), \overline{B}_{\rho_N}(y)) \le 2H \sup_{i \in T} R_I(i)\ \mathrm{len}_D(P), \tag{36.2}$$

$$D(\overline{B}_{\rho_N}(p), \overline{B}_{\rho_N}(q)) \le 4H \big(\inf_{i \in T} R_I(i)\big)^{-1} X_I(P'). \tag{36.3}$$

Here P is a reference path and P' a graph walk, both with range in the fixed buffer; $X_I(P')$ denotes its normalized edge count. A walk of zero edges is handled by the no-traversal case of Lemma 36.1.

Lemma 36.2 (Bounded ratios and a polynomial coarse comparison)**.** *Fix $0 < \kappa < 1/(2\log 2)$ and $N = \lfloor \kappa \log I \rfloor$. There is a fixed $M_0 > 1$ such that, for all sufficiently large I, fewer than one quarter of the labels in $\{N, \dots, 2N\}$ have $R_I(i) \le M_0^{-1}$, and fewer than one quarter have $R_I(i) \ge M_0$. Consequently there are constants $C_0 < \infty$, $\beta_0, \eta_0 > 0$ such that, except with probability $O_K(I^{-\beta_0})$,*

$$X_I(\overline{B}_{I^{-\eta_0}}(x), \overline{B}_{I^{-\eta_0}}(y)) \le C_0 D(x, y), \tag{36.4}$$

$$D(\overline{B}_{I^{-\eta_0}}(p), \overline{B}_{I^{-\eta_0}}(q)) \le C_0 X_I(p, q), \tag{36.5}$$

simultaneously whenever the corresponding source geodesic is contained in K. In (36.5) the endpoints are nuclei. Fixed enlargements of the compact are allowed.

Proof. Take $M_0 = 16H$, increasing it if the fixed geometric conventions require it. Suppose first that at least one quarter of the labels have $R_I(i) \le M_0^{-1}$ along some deterministic $I_n \to \infty$. Use exactly that deterministic set in Lemma 35.1. Its unit intensities satisfy $\min_i J_i \ge I^{1-2\kappa\log 2} \to \infty$, so the local tests apply. Equation (36.2), applied to a reference geodesic between fixed distinct points x, y, gives a connection costing at most $D(x, y)/8$ between ρ_N-perturbed endpoints. A deterministic compact contains this reference geodesic with probability arbitrarily close to one by properness. The perturbed endpoints can be corrected at cost $o_{\mathbb{P}}(1)$ by Proposition 33.3 and continuity of D. This contradicts $X_I(x, y) \to D(x, y) > 0$ in probability.

If at least one quarter of the labels instead satisfy $R_I(i) \ge M_0$, (36.3) gives a reference connection costing at most one quarter of a shortest graph path's normalized cost. Such graph paths between the selected cells of fixed x, y lie in a deterministic large compact with probability arbitrarily close to one. Indeed $X_I(x, y)$ is tight; choose a sphere whose reference exit distance from x exceeds this tight bound with high probability. A first-exit prefix of an escaping graph

geodesic contradicts local uniform convergence on that compact, using the vanishing cell mesh to locate its exit endpoint. Thus $D(x,y) \le X_I(x,y)/4 + o_{\mathbb{P}}(1)$, again a contradiction.

Failure of either eventual deterministic assertion would provide such a sequence I_n. At least half the labels therefore lie in (M_0^{-1}, M_0). Apply coverage to those labels and use (36.2)–(36.3). The cost factor is at most $4HM_0$. The coverage failure is $O_K(e^{-cN}) = O_K(I^{-c\kappa})$. Replace Ce^{-uN} by $I^{-\eta_0}$ with any $\eta_0 < u\kappa$ to absorb fixed constants. This proves the lemma. The argument did not assume a convergence rate or a bound on $\ell(I)/\ell(I^t)$. □

36.2 Improving a comparison constant

We record the deterministic improvement with buffers and graph cuts made explicit. Its joining estimate uses a connected annulus, so it applies in every $d \ge 2$.

Lemma 36.3 (Improvement on a good annular cover)**.** *Let S_0 be either D or X_I on its nuclei, and let T_0 be the other distance. Suppose the baseline bound*

$$T_0(\overline{B}_\rho(x), \overline{B}_\rho(y)) \le CS_0(x,y)$$

holds for all source geodesic subsegments in a fixed buffer. Suppose that buffer is covered by central holes of balls of radii in $[r_-, r_+]$. On every such ball, centered at z with radius r, put $M_B = \{.70r \le |x-z| \le 1.10r\}$ and assume the local tests below hold, with a fixed H_1 and a factor $R_B \le R$:

(i) Every source geodesic crossing in the middle annulus, from the inner strip $[.74r, .76r]$ to the outer strip $[r, 1.01r]$, has a target follow of cost at most $(1+\varepsilon)R_B$ times its source length.

(ii) Let $0 < q_B < \infty$ be the source internal distance between these strips in $A_{r/2,2r}(z)$. The source internal diameter of the inner strip is at most $H_1 q_B$.

(iii) For every point of either strip and every point within Euclidean distance ρ, the target joining cost is at most $\varepsilon R_B q_B$.

(iv) Local filling makes every whole graph cell meeting M_B have Euclidean diameter at most $r/100$, and every edge whose straight segment meets the testing strips have length at most $r/100$. The associated readouts are determined inside a fixed larger annulus.

Then the endpoint tolerance can be enlarged to $4r_+$ and the comparison constant replaced by

$$\Phi(C,R,\varepsilon) = \vartheta \max\{C,R\} + (1-\vartheta)R + 3\varepsilon R, \qquad \vartheta = \frac{H_1}{H_1+1} < 1. \tag{36.6}$$

The constant H_1 can be fixed at the unit failure threshold independently of ε and of the scales.

Proof. Let P be a source geodesic of length L. For a graph geodesic use vertex times, each edge having its normalized cost. Starting at a current point, choose a covering ball with that point in its central half-ball and run to its first exit. For graph paths the exit is the first vertex outside the ball. Continue until the endpoint is reached, recording that endpoint as the last stopping point even when the final block is incomplete. Each complete block has Euclidean displacement at most $2r_+$; the local edge test controls its overshoot. There are finitely many blocks.

Discard an initial and final collection of blocks so that the retained endpoints lie within $4r_+$ of the original endpoints and, for each retained block, both original endpoints lie outside the

selected ball. More explicitly, take the block just after the last block-start within $2r_+$ of the initial endpoint, and stop at the first subsequent block-start within $2r_+$ of the terminal endpoint. If the two endpoints are within $8r_+$, the two output balls already intersect. Otherwise the indicated retained interval exists, and its preceding block is complete.

In a retained block $[t_j, t_{j+1}]$, let s_j be the last visit to radius $3r/4$ before the exit. For a graph path use the last vertex inside that radius. The edge bound places it in the inner strip and the exit vertex in the outer strip; the intervening polygonal path lies in the fixed middle annulus. Write $g_j = t_{j+1} - s_j$ and $b_j = s_j - t_j$. Then $q_B \le g_j$. Because the original starting point is outside the ball, the source path has an earlier inward visit $u_j \le t_j$ to the inner strip. Geodesicity and the intrinsic joining bound give

$$b_j \le s_j - u_j = S_0(P(u_j), P(s_j)) \le H_1 q_B \le H_1 g_j.$$

Thus $b_j \le \vartheta(b_j + g_j)$. This uses an internal strip diameter; it makes no assertion that two unrelated paths intersect.

Apply the baseline comparison to the bad subsegment $[t_j, s_j]$. Its terminal endpoint is in the current inner strip and its initial endpoint is in the preceding ball's outer strip. Repairing both endpoint balls by (iii), and using (i) on the good subsegment, gives

$$T_0(P(t_j), P(t_{j+1})) \le C b_j + (1+\varepsilon) R g_j + \varepsilon R(q_{j-1} + q_j).$$

The complete source blocks are disjoint and each q_j is at most the length of its block. Hence the total endpoint-repair cost is at most $2\varepsilon RL$. With $B = \sum b_j$, $G = \sum g_j$ and $B + G \le L$, the total is at most

$$R(B+G) + (\max\{C, R\} - R)B + 3\varepsilon RL \le \Phi(C, R, \varepsilon) L.$$

The coefficient multiplying B is nonnegative, so this step remains valid when $C < R$. All graph cuts were vertices; there is no sum of uncontrolled rounding errors.

Lemma 34.1 supplies (i)–(iii) under the unit law. The source strip diameter divided by its crossing distance has a fixed high-probability cutoff H_1 for both source metrics. Thus H_1 is chosen before ε. If $\rho/r_- \to 0$, the endpoint-repair condition has probability tending to one. Local filling supplies (iv). □

36.3 Selecting enough ratios close to one

Lemma 36.4 (Polynomial-probability near-one comparison)**.** *For every $\delta \in (0,1)$ there are $\beta, \eta > 0$ such that, for every fixed compact buffer, except with probability $O_K(I^{-\beta})$,*

$$X_I(\overline{B}_{I^{-\eta}}(x), \overline{B}_{I^{-\eta}}(y)) \le (1+\delta) D(x,y), \tag{36.7}$$

$$D(\overline{B}_{I^{-\eta}}(p), \overline{B}_{I^{-\eta}}(q)) \le (1+\delta) X_I(p,q), \tag{36.8}$$

simultaneously for all source geodesics contained in the buffer. The second line concerns graph nuclei. The exponents may depend on δ.

Proof. Start with Lemma 36.2, using a fixed κ_0 and endpoint exponent $\eta_0 = u\kappa_0/2$. Choose a decreasing sequence $\kappa_j > 0$ so rapidly that

$$v\kappa_j < \tfrac{1}{4} u \kappa_{j-1} \quad (j \ge 1). \tag{36.9}$$

The jth label window is $W_j(I) = \{N_j, \dots, 2N_j\}$ with $N_j = \lfloor \kappa_j \log I \rfloor$. All effective intensities tend to infinity. After window j enlarge the output endpoint tolerance, if necessary, to $e_j(I) = I^{-u\kappa_j/2}$; its actual radius $4r_+$ is smaller than this for all large I. Before window j, use e_{j-1} as the baseline tolerance, also when some intervening windows were skipped. For every label and physical radius in that window impose the same unit repair radius

$$t_{I,j} = C_* e_{j-1}(I) e^{vN_j} \le I^{-u\kappa_{j-1}/4} \longrightarrow 0.$$

Here C_* absorbs the fixed lower-radius factor in Lemma 35.1; the displayed bound holds for all large I by the strict inequality in (36.9). This dominates the physical baseline tolerance divided by the radius. Equation (34.1), with $\min_{i \in W_j} J_i \to \infty$, shows that the unit repair-test failure probability tends to zero uniformly over the labels. Thus each label has one fixed unit event to which Lemma 35.1 applies; no event is being chosen after observing its radius or its coarse history. There are only finitely many windows in each application.

Choose $t \in (0, 1/2)$ and then $\varepsilon > 0$ sufficiently small that

$$F := 1 + \frac{3\varepsilon}{1-\vartheta} \quad \text{satisfies} \quad F(1-t) < 1 - t/2, \qquad \frac{F}{1-t} < 1 + \delta/2. \tag{36.10}$$

Choose an integer m so large that iterating (36.6) m times with $R \le 1-t$, starting from C_0, produces a constant less than $1 - t/4$, and also $\vartheta^m \max\{C_0, (1-t)^{-1}\} < \delta/2$. One explicit majorant for the first requirement is

$$c_k = \vartheta^k \max\{C_0, 1-t\} + F(1-t)(1-\vartheta^k),$$

which is at least $1-t$ and obeys the required recurrence bound.

Call a window low if at least one third of its labels have $R_I(i) \le 1-t$. Call it high if at least one third have $R_I(i)^{-1} \le 1-t$. Among the first $3m$ windows there are eventually fewer than m low windows and fewer than m high windows. Indeed, otherwise pass to a deterministic sequence of intensities and a fixed choice of m offending windows. Apply coverage and contraction in those windows in increasing order. For low windows use source D and target X_I; for high windows interchange the two. The comparison factor becomes less than $1 - t/4$ with probability tending to one. At fixed distinct endpoints this contradicts Proposition 33.3, exactly as in Lemma 36.2. The reference geodesic and the graph geodesic are confined with arbitrarily high probability, and the polynomially vanishing endpoint balls contribute $o_{\mathbb{P}}(1)$. This contradiction uses no rate for the original convergence.

Consequently at least m of those windows have at least one third of their labels in the interval

$$1 - t < R_I(i) < (1-t)^{-1}. \tag{36.11}$$

These are deterministic label sets for each I. Use them for both directions of the improvement. Their factors R and R^{-1} are at most $(1-t)^{-1}$. The constants are bounded by the recurrence with this R; its limiting majorant is $F/(1-t)$. The two choices in (36.10) and the choice of m make the final constant less than $1 + \delta$.

Each of the finitely many coverage events fails with probability at most $C_K e^{-cN_j} = O_K(I^{-c\kappa_j})$. Their union and the baseline event therefore have a positive polynomial failure exponent. Their endpoint radii are bounded by $e_{3m}(I) = I^{-u\kappa_{3m}/2}$ for all sufficiently large I; take any $\eta < u\kappa_{3m}/2$. This proves (36.7)–(36.8). Neither β nor η is asserted to stay bounded away from zero as $\delta \downarrow 0$. □

37 Interpolation for the marked Poisson process

The graph distance is not monotone in the intensity. The following argument controls all intermediate graphs simultaneously by bounding clusters of newly added sites.

Lemma 37.1 (Sparse birth clusters on a polynomial grid)**.** *Fix $p > 0$, put $s_n = n^p$, $b_n = (n+1)^p$, and let $E_n = \mathcal{P}_{b_n} \setminus \mathcal{P}_{s_n}$. There are constants $0 < \kappa < \beta < \alpha$, $q > 0$ and an integer m, such that, with*

$$r_n = s_n^{-\alpha+\kappa}, \qquad R_n = s_n^{-\alpha+\beta}, \qquad M_n = s_n^q,$$

almost surely eventually on every fixed compact buffer:

(i) all cells of every actual intermediate cloud have whole D-diameter at most r_n;

(ii) every D-ball of radius $10r_n$ has at most M_n nuclei of $\mathcal{P}_{b_n}$;

(iii) every component of the graph on E_n in that buffer with edges between sites at D-distance at most R_n has fewer than m sites.

The parameters can be chosen so that

$$M_n r_n / R_n \longrightarrow 0, \qquad a(s_n) M_n \longrightarrow 0. \tag{37.1}$$

Proof. Conditional on h, E_n is Poisson of intensity $\lambda_n \mu$, where $\lambda_n = b_n - s_n \le C_p s_n^{1-1/p}$. Choose

$$0 < \beta < \min\{\alpha/4, 1/(4p\Delta)\}, \qquad 0 < \eta < 1/(4p\alpha).$$

By (2.2), eventually

$$\lambda_n \sup_{z \in V} \mu(B_D(z, 2R_n)) \le C s_n^{-1/p+\beta\Delta+\alpha\eta-\beta\eta} \le C s_n^{-1/(2p)}.$$

The conditional Mecke formula and a sum over labeled trees on m vertices give

$$\begin{aligned} &\mathbb{P}[\text{some birth component in } V \text{ has at least } m \text{ vertices} \mid h] \\ &\quad \le C_m \lambda_n \mu(V) \left(\lambda_n \sup_{z \in V} \mu(B_D(z, 2R_n)) \right)^{m-1} \le C_{h,m,p} n^{p-1-(m-1)/2}. \end{aligned}$$

A component with at least m vertices contains a tree on exactly m distinct vertices, which proves the first bound without a degree assumption. The radius $2R_n$ in the open balls includes edges of length exactly R_n without any assumption about the mass of metric spheres. Choose $m > 2p + 1$. Conditional Borel–Cantelli proves (iii).

Choose $\kappa, q > 0$ with

$$\kappa\Delta < q, \qquad q + \kappa < \beta, \qquad q < \alpha.$$

Assertion (i) follows from the arbitrary-slab consequence of (15.4) recorded at the start of Section 34, since $b_n \le 2s_n$ eventually. The conditional Poisson grid proof of (34.4), now with upper cloud $\mathcal{P}_{b_n}$, threshold s_n^q, and sufficiently small volume loss, proves (ii). Its tail is summable even along $s_n = n^p$. These choices imply (37.1). Intersect over integer buffers. Independence of successive slabs has not been used. □

Lemma 37.2 (Deterministic comparison throughout a slab)**.** *On the events of Lemma 37.1, write $s = s_n$, $b = b_n$, $r = r_n$, $R = R_n$, $M = M_n$. A simple N-edge path of any intermediate cloud $\mathcal{P}_I$, $s \le I \le b$, confined to a fixed buffer, has a lower-cloud follow of at most*

$$(1 + 2mMr/R)N + (m + 2)M \tag{37.2}$$

edges. Conversely a simple N-edge lower-cloud path has a follow in every intermediate cloud with at most

$$(1 + 4mr/R)N + m + 2M \tag{37.3}$$

edges. The endpoints are the selected containing cells of the same Euclidean queries. These follows lie in an arbitrarily small fixed Euclidean enlargement of the original buffer for all large n.

Proof. Use a slightly larger buffer for all nuclei and birth components. Adjacent current nuclei have D-distance at most $2r$, since their cells touch. List the birth vertices of a simple current path in path order, and divide the list into maximal successive runs belonging to one R-component of the birth set. Each run has fewer than m vertices. Between consecutive runs the path travels D-distance greater than R, so it uses more than $R/(2r)$ edges. Thus the number L of birth vertices on the path satisfies

$$L \le m(2rN/R + 1). \tag{37.4}$$

This is a deterministic statement about every path, including paths selected after the whole slab is observed.

Choose a continuous trace through the source cells, through their nuclei as interior waypoints, crossing chosen facets at relative-interior points. For an old nucleus, its current cell is contained in its lower-cloud cell. For a birth cell with nucleus z, every lower cell meeting the trace portion in it has nucleus w with $D(z, w) \le 2r$, so there are at most M such cells. The complete family of lower cells meeting the trace is facet connected. The old portions contribute at most $N + 1$ cells; the birth portions contribute at most ML. At the two endpoints, all extra incident nuclei lie in a D-ball of radius at most $2r$, so at most $2M$ additional cells handle every tie choice. This gives $N + ML + 2M$ edges and hence (37.2).

For the other direction, take a simple lower path with nuclei $v_0, \dots, v_N$ and the same type of trace. A current cell with an old nucleus can meet its interior only in the lower cell of that nucleus. Thus at most $N + 1 + 2M$ old current cells are needed, including endpoint ties. For every distinct birth cell meeting the trace, assign its nucleus z to an index i whose lower cell contains a contact point. Then $D(z, v_i) \le 2r$. Order these distinct birth nuclei by assigned index. If two successive entries z, w in different birth components are assigned to $i \le j$, then

$$R < D(z, w) \le 4r + 2r(j - i).$$

For large n, $R \ge 8r$, so $j - i > R/(4r)$. Again each run within a component has fewer than m sites. Hence at most $m(4rN/R + 1)$ birth cells meet the trace. Facet connectivity gives (37.3).

These contact counts remain valid if the trace meets an additional target face: all target cells meeting a birth portion are included in its M bound. A source-old cell interior has that old nucleus uniquely nearest among the lower sites as well. The only extra old contacts at source facet transitions are the two source neighbors, with endpoint ties already counted. Finally the whole Euclidean cell mesh tends to zero uniformly over the slab. Every target cell meets the trace, so the resulting routes lie in the asserted enlargement. □

38 Proof of Theorem 1.1(i)

Fix $\delta \in (0,1)$ and obtain β, η from Lemma 36.4. Choose an integer $p > 1/\beta$ and set $s_n = n^p$. Borel–Cantelli gives the near-one comparisons at all sufficiently large s_n, simultaneously on each integer-radius buffer. Intersect with the probability-one events of Lemmas 34.2 and 37.1, and the reference estimates of Section 2 and the cell estimates of Section 15. All statements in what follows are pathwise on this event.

Correcting endpoints and confining grid geodesics. Fix a compact K. Properness and continuity permit a finite integer radius containing all reference geodesics between pairs of points of K. Indeed these geodesics lie in a bounded reference ball about any one point of K. Apply (36.7) to them and correct the two endpoint balls using Lemma 34.2. Uniformly on K^2,

$$X_{s_n}(x,y) \le (1+\delta)D(x,y) + o(1). \tag{38.1}$$

Choose a larger integer radius L with

$$D(K, \partial B_L) > 10(\operatorname{diam}_D K + 1).$$

If a graph geodesic between selected cells of points of K exits B_L, stop at its first exit nucleus (or at the next nucleus of its continuous cell trace). Its prefix is still a graph geodesic, and lies in B_{L+1} for large n by the vanishing whole cell mesh. The lower comparison (36.8) applied to this prefix, together with uniform continuity of D, gives normalized prefix cost at least

$$(1+\delta)^{-1} D(K, \partial B_L) - o(1).$$

This exceeds the upper bound in (38.1). Thus all grid geodesics joining K are eventually confined to one fixed enlarged buffer. Apply (36.8) to each complete graph geodesic. Its endpoint nuclei approach their queries uniformly by the Euclidean mesh, and the reference endpoint balls have vanishing diameter. Hence

$$D(x,y) \le (1+\delta)X_{s_n}(x,y) + o(1) \quad \text{uniformly on } K^2. \tag{38.2}$$

The argument applies to every compact on the same event. In particular the grid distances are locally bounded, and their uniform error on K has limsup at most $\delta \operatorname{diam}_D K$.

All real intensities. Use Lemma 37.2 for the grid exponent p. Following lower-grid geodesics by (37.3), and using (37.1), gives a uniform bound for $a(s_n)d_I(x,y)$, $x, y \in K$, throughout $s_n \le I \le s_{n+1}$. Every such route is confined to a fixed enlarged buffer.

Intermediate geodesics are confined too. Choose an integer sphere whose reference exit distance from K is sufficiently larger than the preceding uniform bound. If an intermediate geodesic exits it, project its first-exit prefix to the lower cloud by (37.2). The projected lower distance is at most the intermediate geodesic's normalized cost plus $o(1)$. But (38.2), now on the compact containing K and the exit sphere, gives a strictly larger lower bound. This is impossible. All the involved compacts are fixed after choosing the sample; the countable intersection over integer buffers has already supplied their estimates.

We may therefore apply both path surgeries to the global shortest paths. Their multiplicative errors tend to zero, their additive errors vanish after multiplication by $a(s_n)$, and the normalized

distances are bounded. It follows that

$$\sup_{s_n \le I \le s_{n+1}} \sup_{x,y \in K} |a(s_n) d_I(x,y) - a(s_n) d_{s_n}(x,y)| \longrightarrow 0. \tag{38.3}$$

The uniform convergence theorem for measurable regularly varying functions [17, Theorem 1.2.1], applied to (1.6), and $s_{n+1}/s_n \to 1$, gives

$$\sup_{s_n \le I \le s_{n+1}} \left| \frac{a(I)}{a(s_n)} - 1 \right| \longrightarrow 0.$$

Combining this with (38.3) and the two grid inequalities yields

$$\limsup_{I \to \infty} \sup_{x,y \in K} |X_I(x,y) - D(x,y)| \le \delta \operatorname{diam}_D K \quad \text{almost surely.} \tag{38.4}$$

The single event and the order of quantifiers. Repeat the argument for $\delta = 1/m$, $m \ge 2$, choosing a separate integer grid exponent p_m when necessary. Take the countable intersection of their probability-one events and of the events for all integer-radius buffers. Equation (38.4) then holds on this one event for every m and every compact K. Sending $m \to \infty$ proves (1.5). Every compact is contained in an integer-radius ball, so the compact quantifier introduces no further exceptional events.

The interpolation estimates were deterministic on each marked slab and held for every actual intermediate cloud simultaneously. Thus this last step concerns all real intensities of the marked Poisson process (1.3), rather than only a countable sequence or a recoupled version. □

References

[1] OpenAI. *Metric-measure limits of subcritical FK and spanning-tree planar maps*. OpenAI Math Release preprint, September 24, 2026. Inspected source revision `4eaee87548250`. The passage-law framework is substantially adapted in Sections 16–31; the planar-map convergence theorem is not used as an all-dimensional convergence input. No formal verification of that preprint is asserted here.

[2] *Higher-dimensional LQG confluence*. Research report, November 2, 2026, 57 pages. Source file `higher_dimensional_lqg_confluence(1).pdf`. Theorem 1.2 and Proposition 3.1 are the reference-metric inputs used for $d \geq 3$. Its contact theorem belongs to the input chain of the strong-confluence construction in [3].

[3] *Higher-dimensional strong confluence: proof and research report*. Research report, November 3, 2026, 32 pages. Source file `higher_dimensional_strong_confluence_progress(1).pdf`. Constructive core coverage is used from Lemma 3.2, Proposition 5.1 and Sections 5.7, 7.2–7.5, with the local escape observation proved in Section 22.

[4] Ewain Gwynne and Jason Miller. Existence and uniqueness of the Liouville quantum gravity metric for $\gamma \in (0, 2)$. *Inventiones Mathematicae* **223** (2021), 213–333. doi:10.1007/s00222-020-00991-6.

[5] Jason Miller and Wei Qian. The geodesics in Liouville quantum gravity are not Schramm–Loewner evolutions. *Probability Theory and Related Fields* **177** (2020), 677–709. doi:10.1007/s00440-019-00949-7.

[6] Morris Ang, Hugo Falconet and Xin Sun. Volume of metric balls in Liouville quantum gravity. *Electronic Journal of Probability* **25** (2020), article 160, 1–50. arXiv:2001.11467v2.

[7] Jian Ding, Ewain Gwynne and Zijie Zhuang. Tightness of exponential metrics for log-correlated Gaussian fields in arbitrary dimension. arXiv:2310.03996, version 5, 2026.

[8] Andres A. Contreras Hip and Zijie Zhuang. Weak exponential metrics for high-dimensional log-correlated Gaussian fields. arXiv:2512.06292v1, December 6, 2025.

[9] Rémi Rhodes and Vincent Vargas. Gaussian multiplicative chaos and applications: a review. *Probability Surveys* **11** (2014), 315–392. arXiv:1305.6221.

[10] Raoul Robert and Vincent Vargas. Gaussian multiplicative chaos revisited. *Annals of Probability* **38** (2010), 605–631. arXiv:0807.1030. Proposition 3.6 supplies the negative moments of reference ball masses.

[11] Manan Bhatia and Konstantinos Kavvadias. Strong confluence of geodesics in Liouville quantum gravity. arXiv:2512.09219, version 2, 2026. This is a planar theorem; it is not invoked as an arbitrary-dimensional result.

[12] Andres A. Contreras Hip and Zijie Zhuang. Bounds on the distance exponent for higher-dimensional Liouville first passage percolation. arXiv:2504.09141, 2025.

[13] Geoffrey R. Grimmett, Alexander E. Holroyd and Gady Kozma. Percolation of finite clusters and infinite surfaces. *Mathematical Proceedings of the Cambridge Philosophical Society* **156** (2014), 263–279. Lemma 3.1 is the plaquette-boundary connectivity input. Author-hosted paper.

[14] *Higher-dimensional exponential metrics: a proposed axiomatic existence and uniqueness proof.* Research manuscript, October 2, 2026, canonical 32-page version, source revision `f1c36687a4f5`. Definition 1.1, Theorems 1.2 and 1.4, Lemmas 2.6 and 2.7, and Lemma 10.1 are the reference inputs used here. Its proposed-proof status is retained. No independent replay of the formalizations reported in [15] is claimed.

[15] Leonardo Bonanno. Research page. Source record consulted November 12–13, 2026. Formalization statements are website reports; no Lean proof object is used as a verified input here.

[16] Charles Devlin VI. Almost Sure Convergence of Liouville First Passage Percolation. *Probability Theory and Related Fields* **193** (2025), 247–288. arXiv:2309.08001. The annular improvement and normalizer-selection strategy is developed in Sections 3–4.

[17] N. H. Bingham, C. M. Goldie and J. L. Teugels. *Regular Variation.* Encyclopedia of Mathematics and its Applications, vol. 27. Cambridge University Press, 1987. doi:10.1017/CBO9780511721434.